



US007106336B1

(12) **United States Patent**
Hutchins

(10) **Patent No.:** **US 7,106,336 B1**
(45) **Date of Patent:** **Sep. 12, 2006**

(54) **METHOD AND SYSTEM FOR DEFERRED EVALUATION OF TRANSFORMS IN GRAPHICS PROCESSORS**

(56) **References Cited**

U.S. PATENT DOCUMENTS

(75) Inventor: **Edward A. Hutchins**, Mountain View, CA (US)

5,801,711 A * 9/1998 Koss et al. 345/441
6,198,488 B1 * 3/2001 Lindholm et al. 345/426
6,259,461 B1 * 7/2001 Brown 345/556

(73) Assignee: **nVidia Corporation**, Santa Clara, CA (US)

* cited by examiner

Primary Examiner—Kee M. Tung
Assistant Examiner—Joni Hsu

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(57) **ABSTRACT**

(21) Appl. No.: **10/846,787**

A method of deferring evaluation of a transform, in accordance with one embodiment of the present invention, includes buffering a plurality of vertex data. The method also includes receiving a draw command, accessing a given vertex data corresponding to the draw command and an associated transform indicator bit. The given vertex data is transformed if the associated indicator bit is cleared. After performing the transform, the vertex data is overwritten with the transformed vertex data and the associated transform indicator bit is set.

(22) Filed: **May 14, 2004**

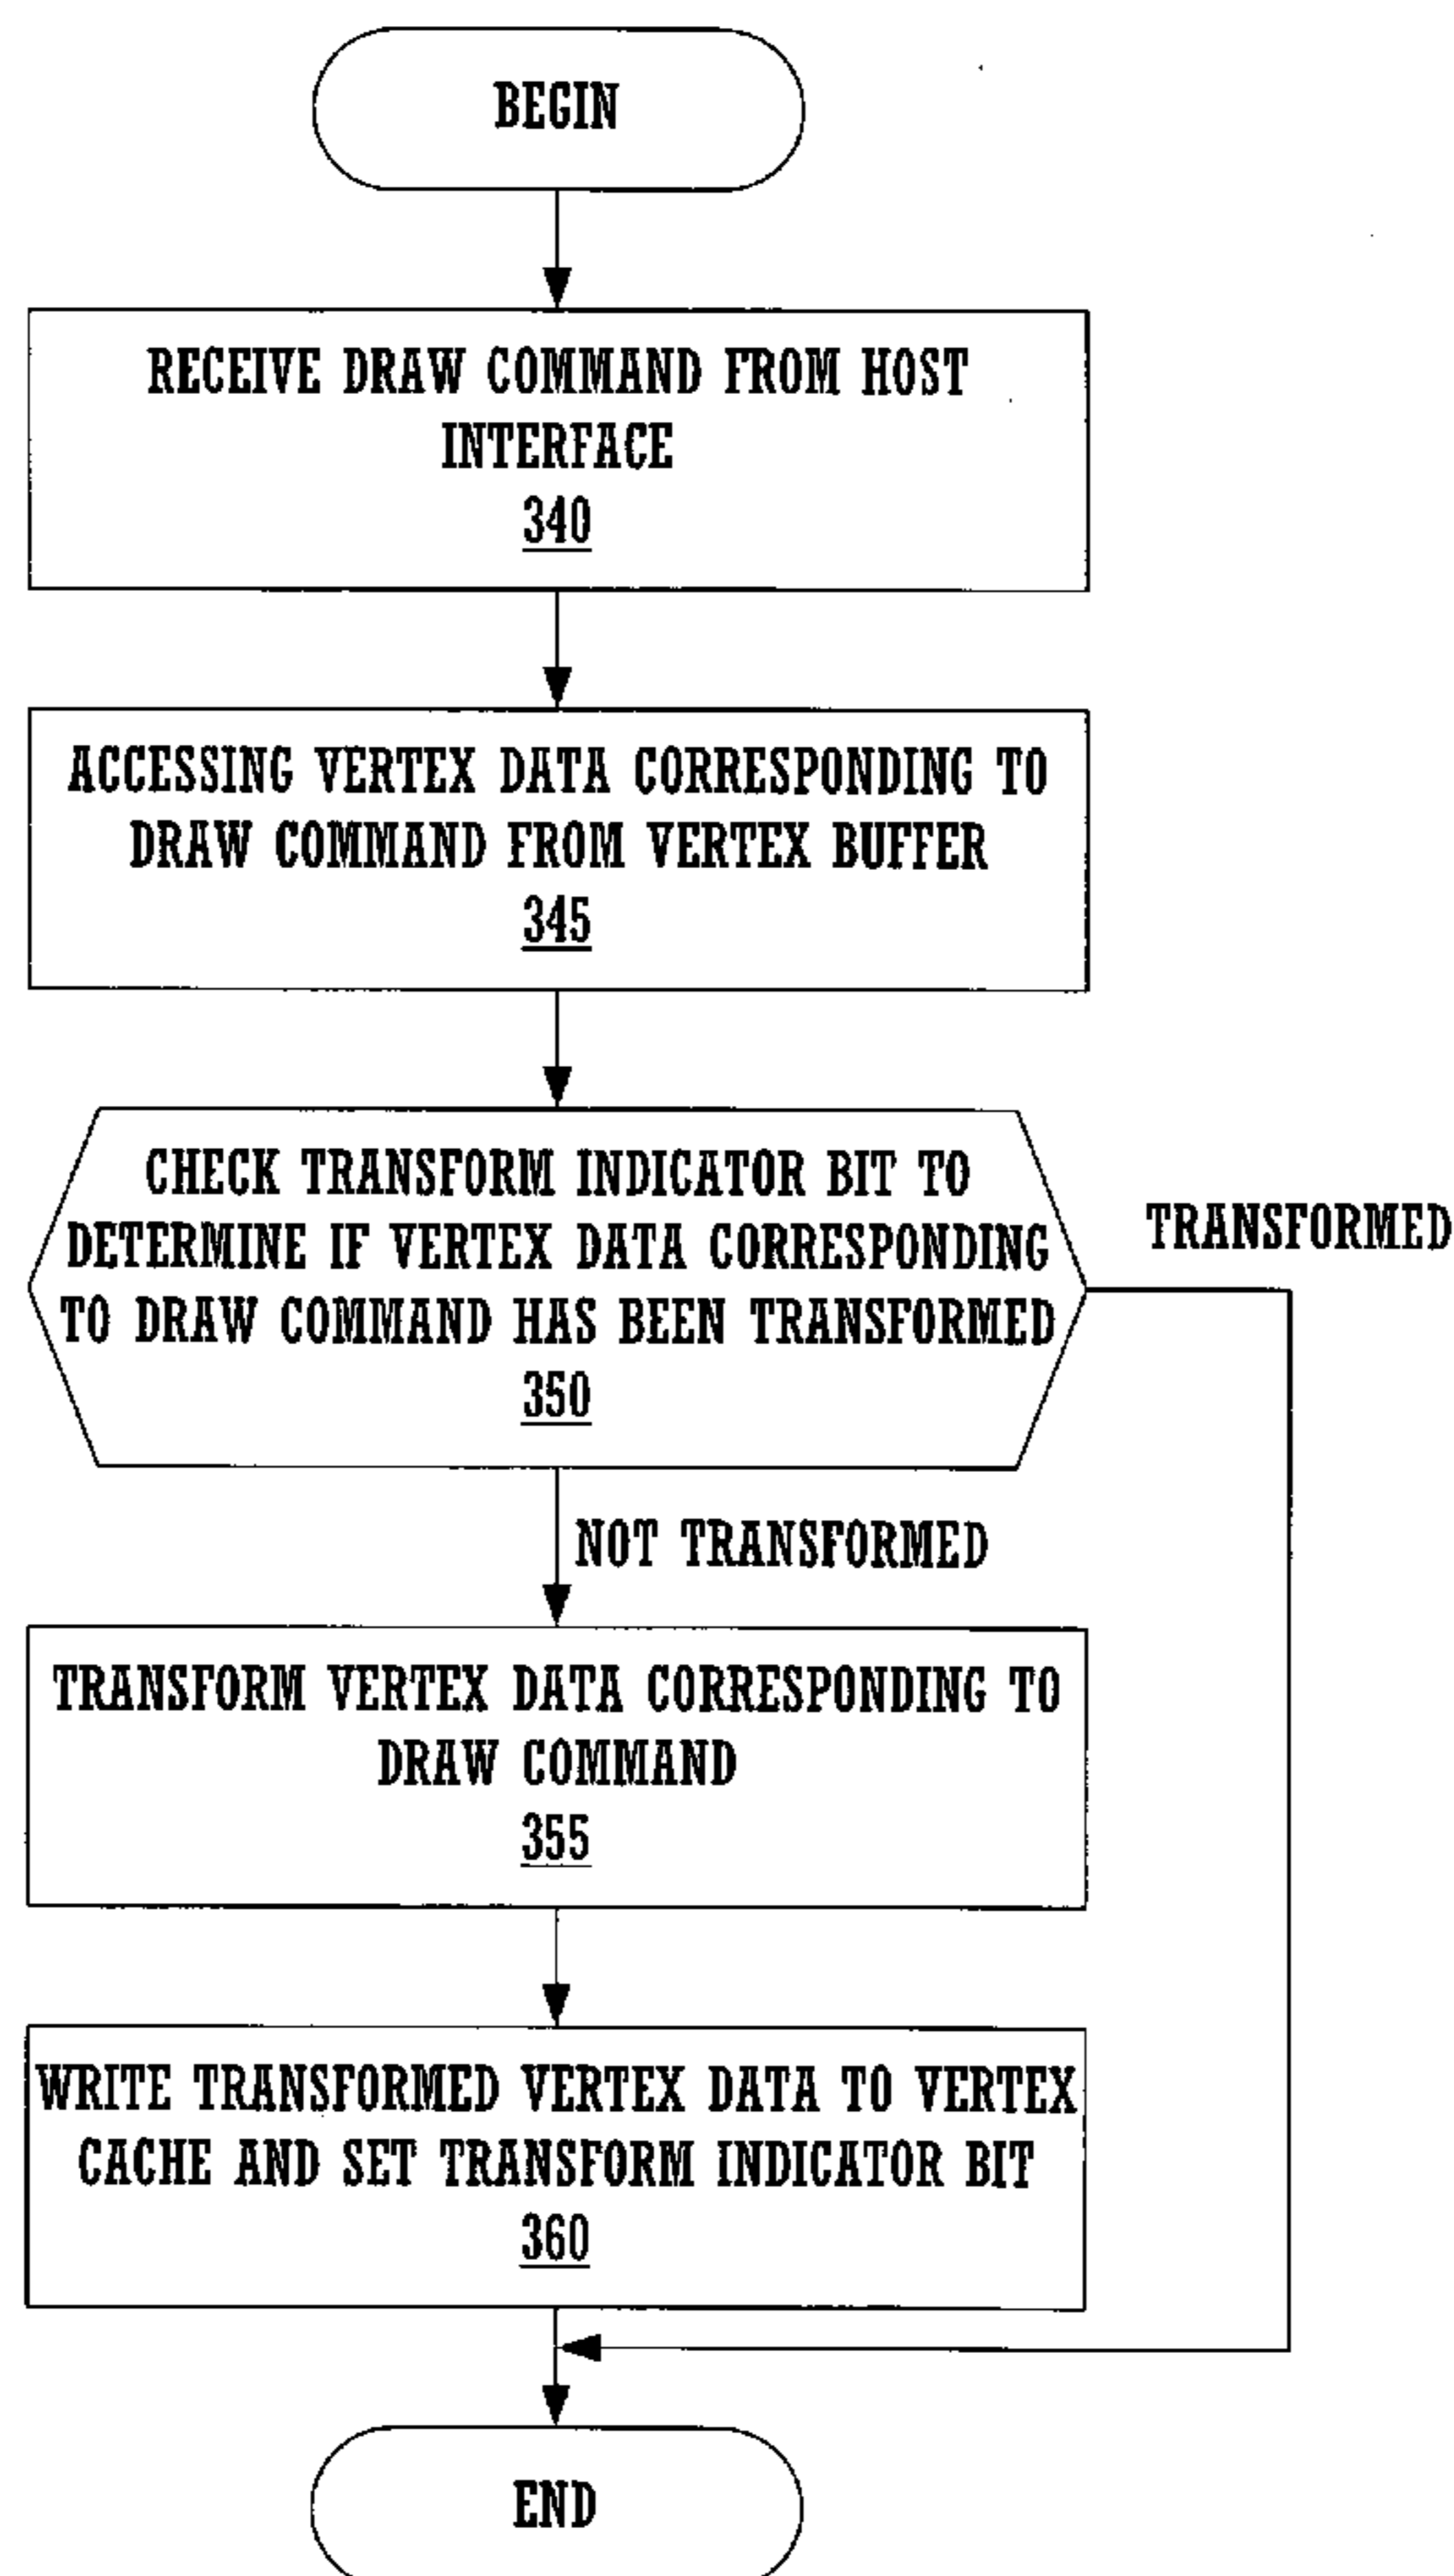
(51) **Int. Cl.**
G06T 1/60 (2006.01)
G06T 15/10 (2006.01)
G09G 5/36 (2006.01)

(52) **U.S. Cl.** **345/530; 345/556; 345/427**

(58) **Field of Classification Search** **345/530, 345/427, 419, 556, 418**

See application file for complete search history.

20 Claims, 9 Drawing Sheets



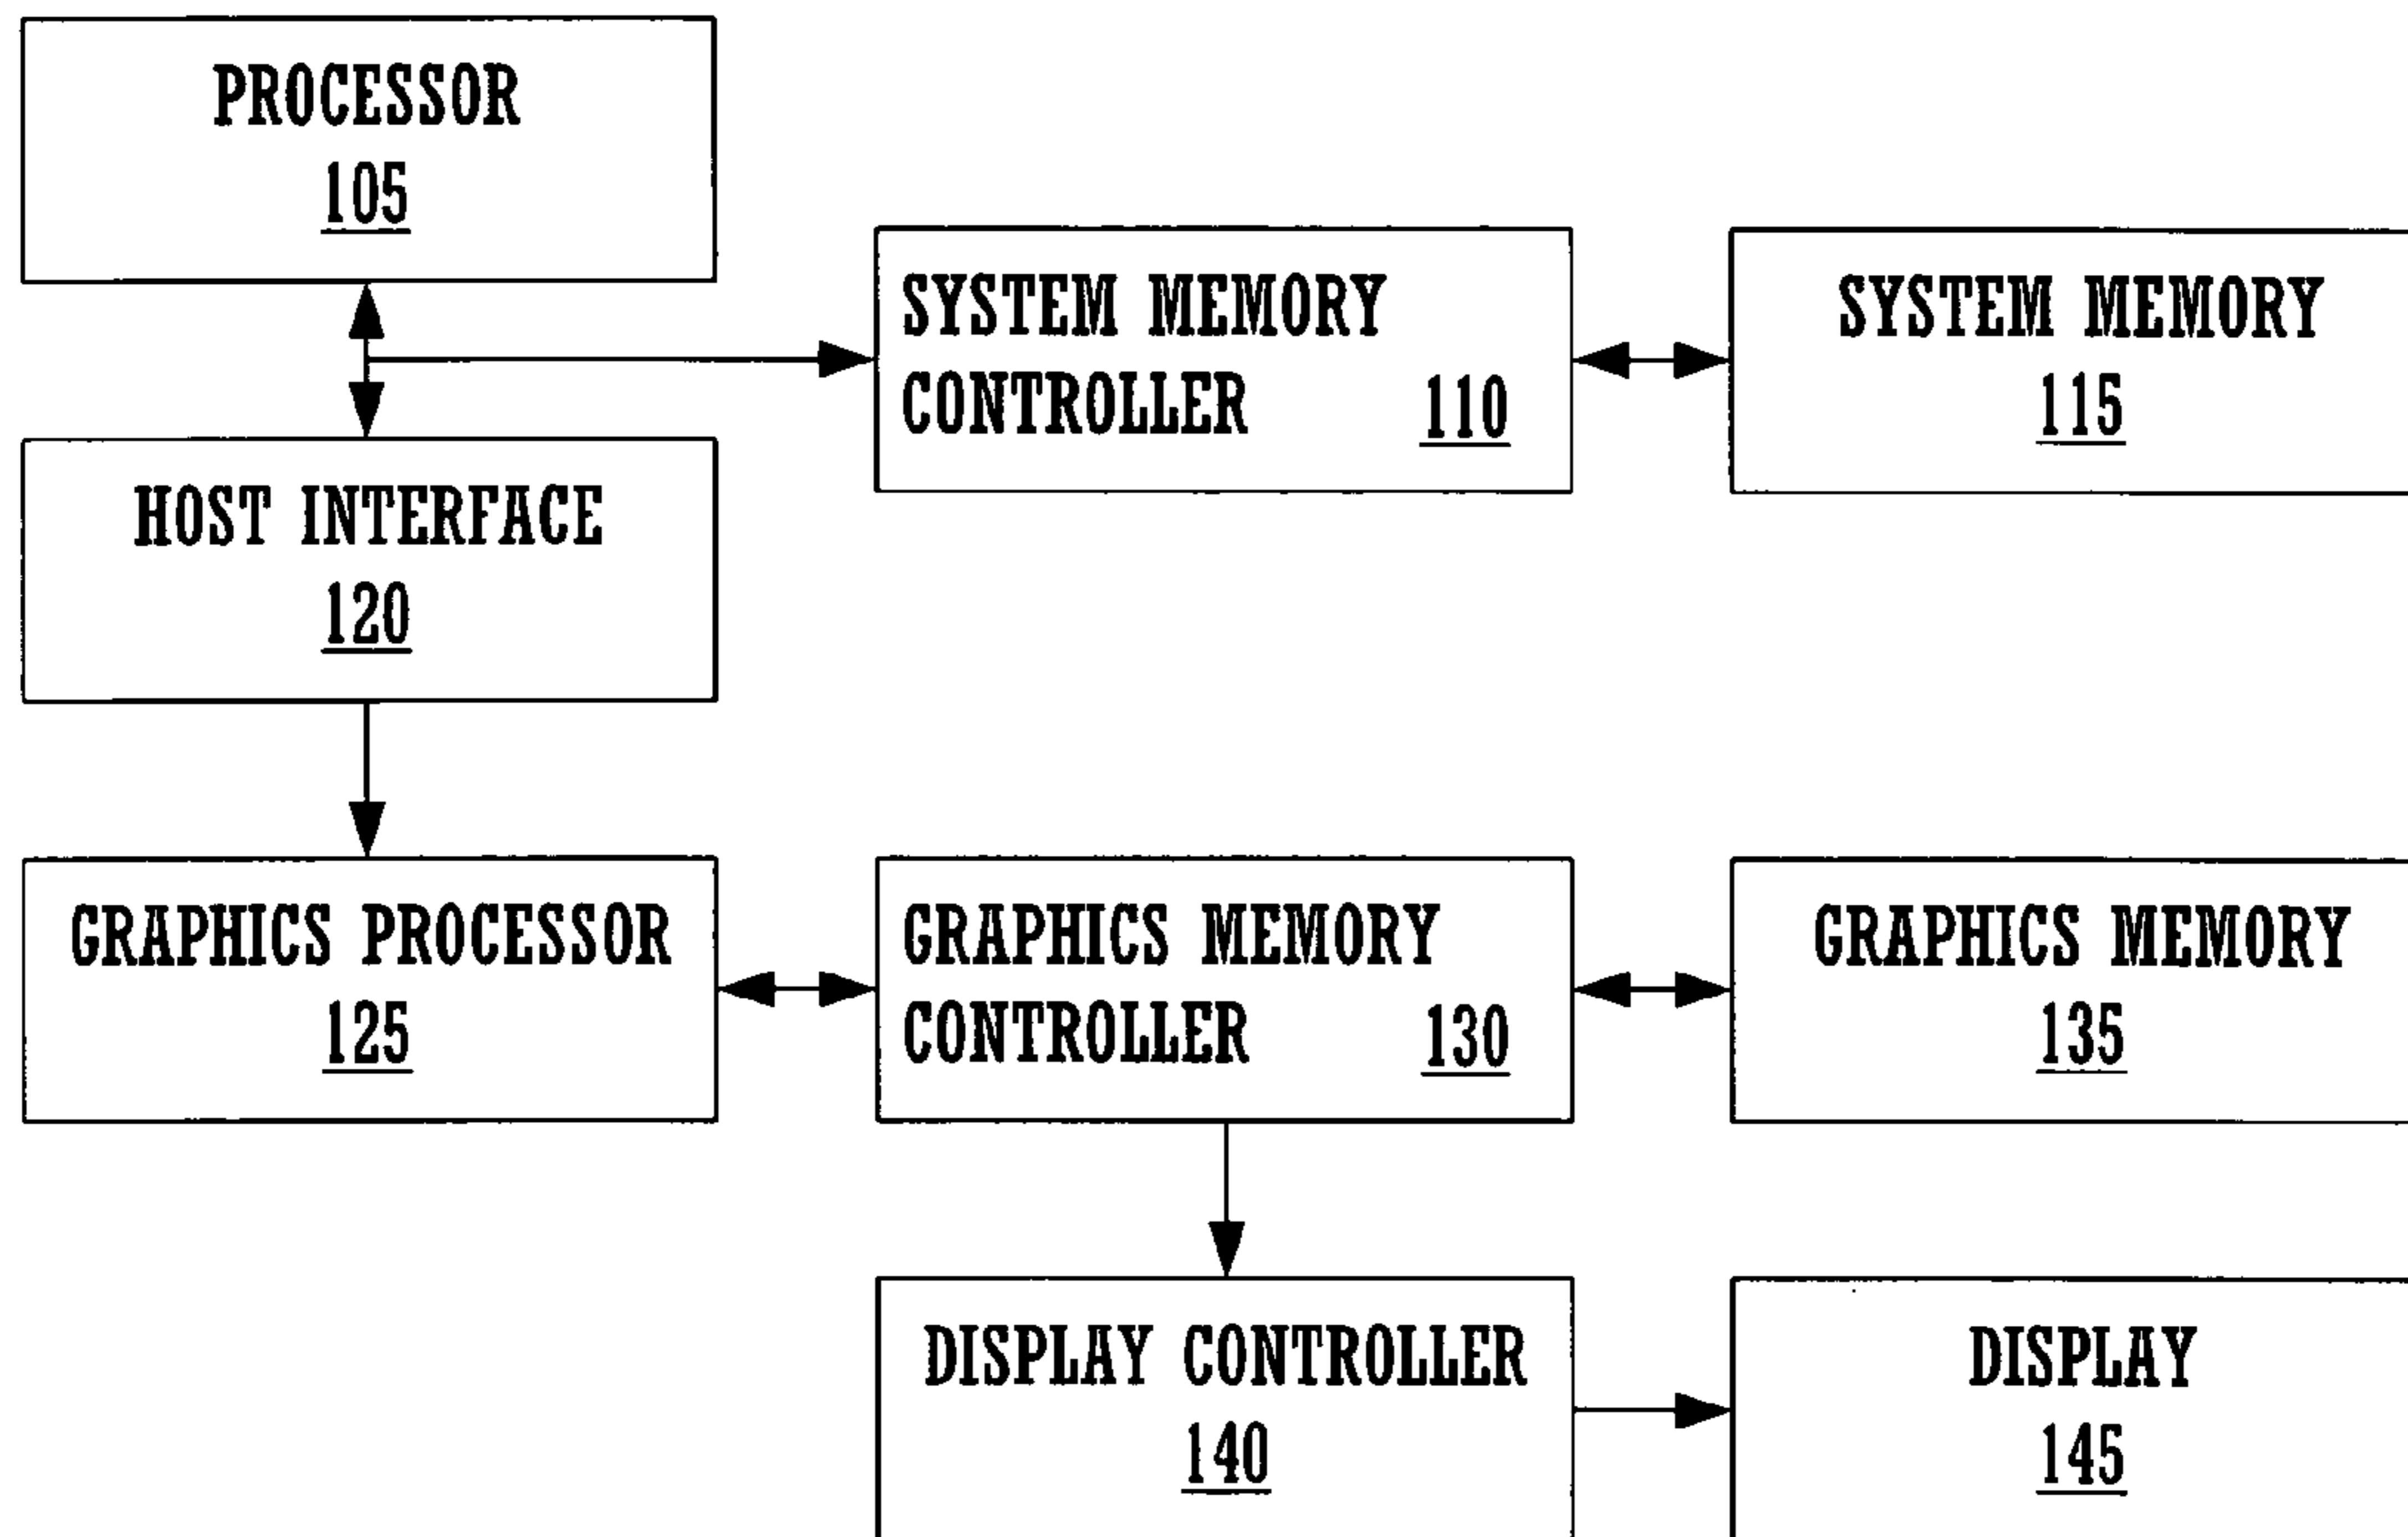


FIGURE 1A

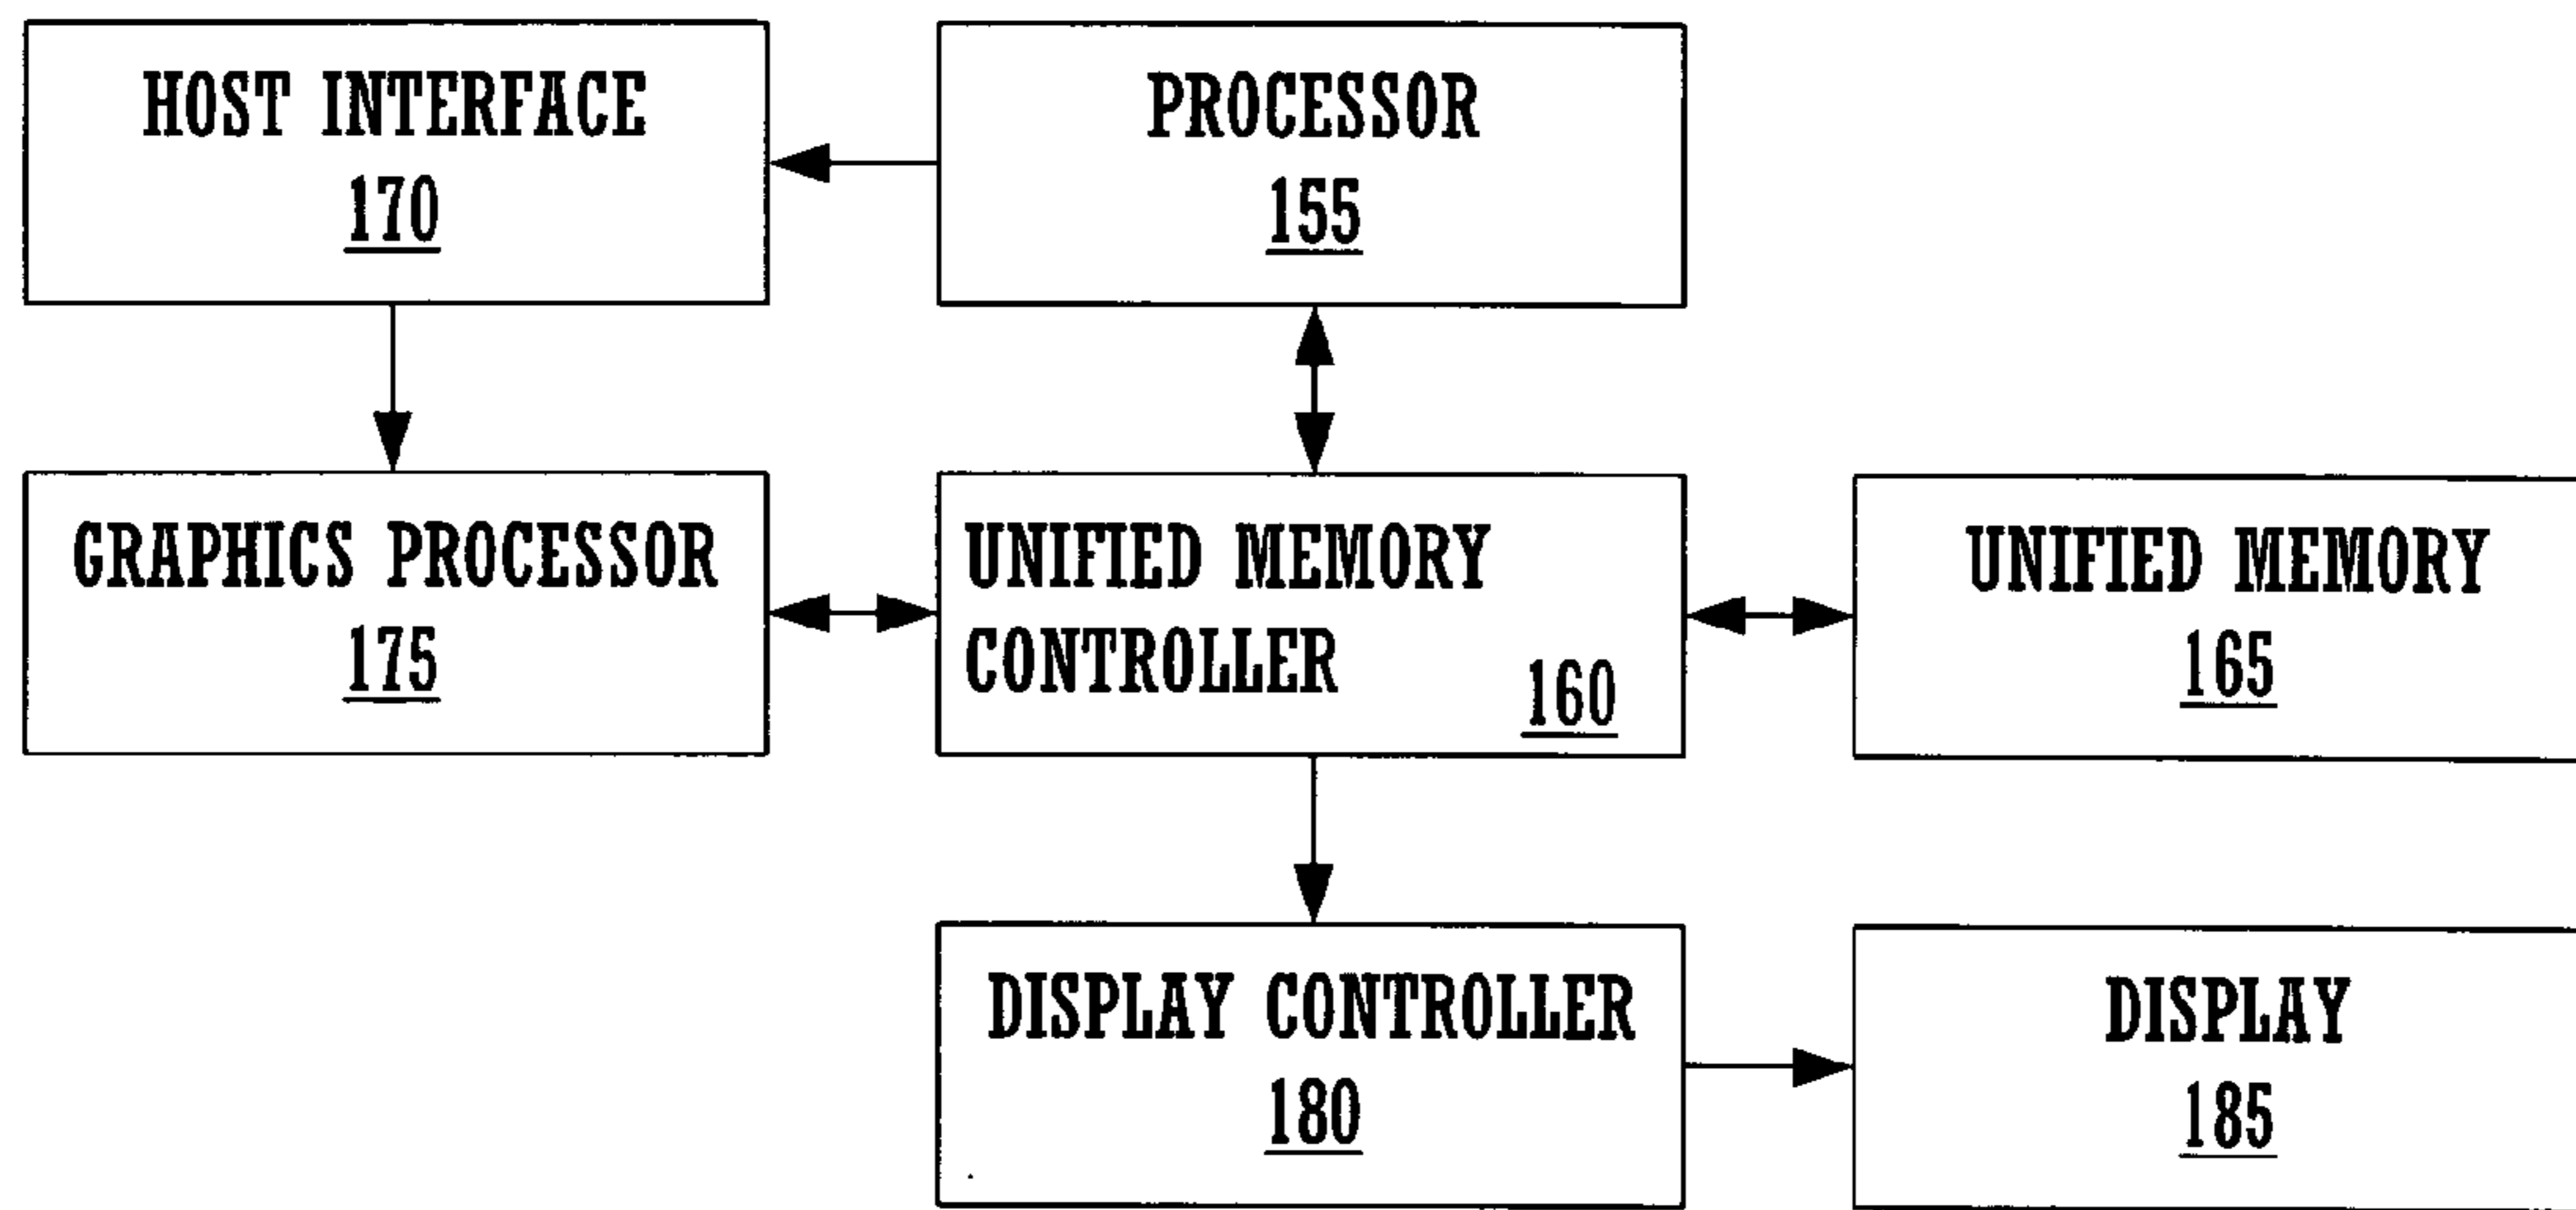


FIGURE 1B

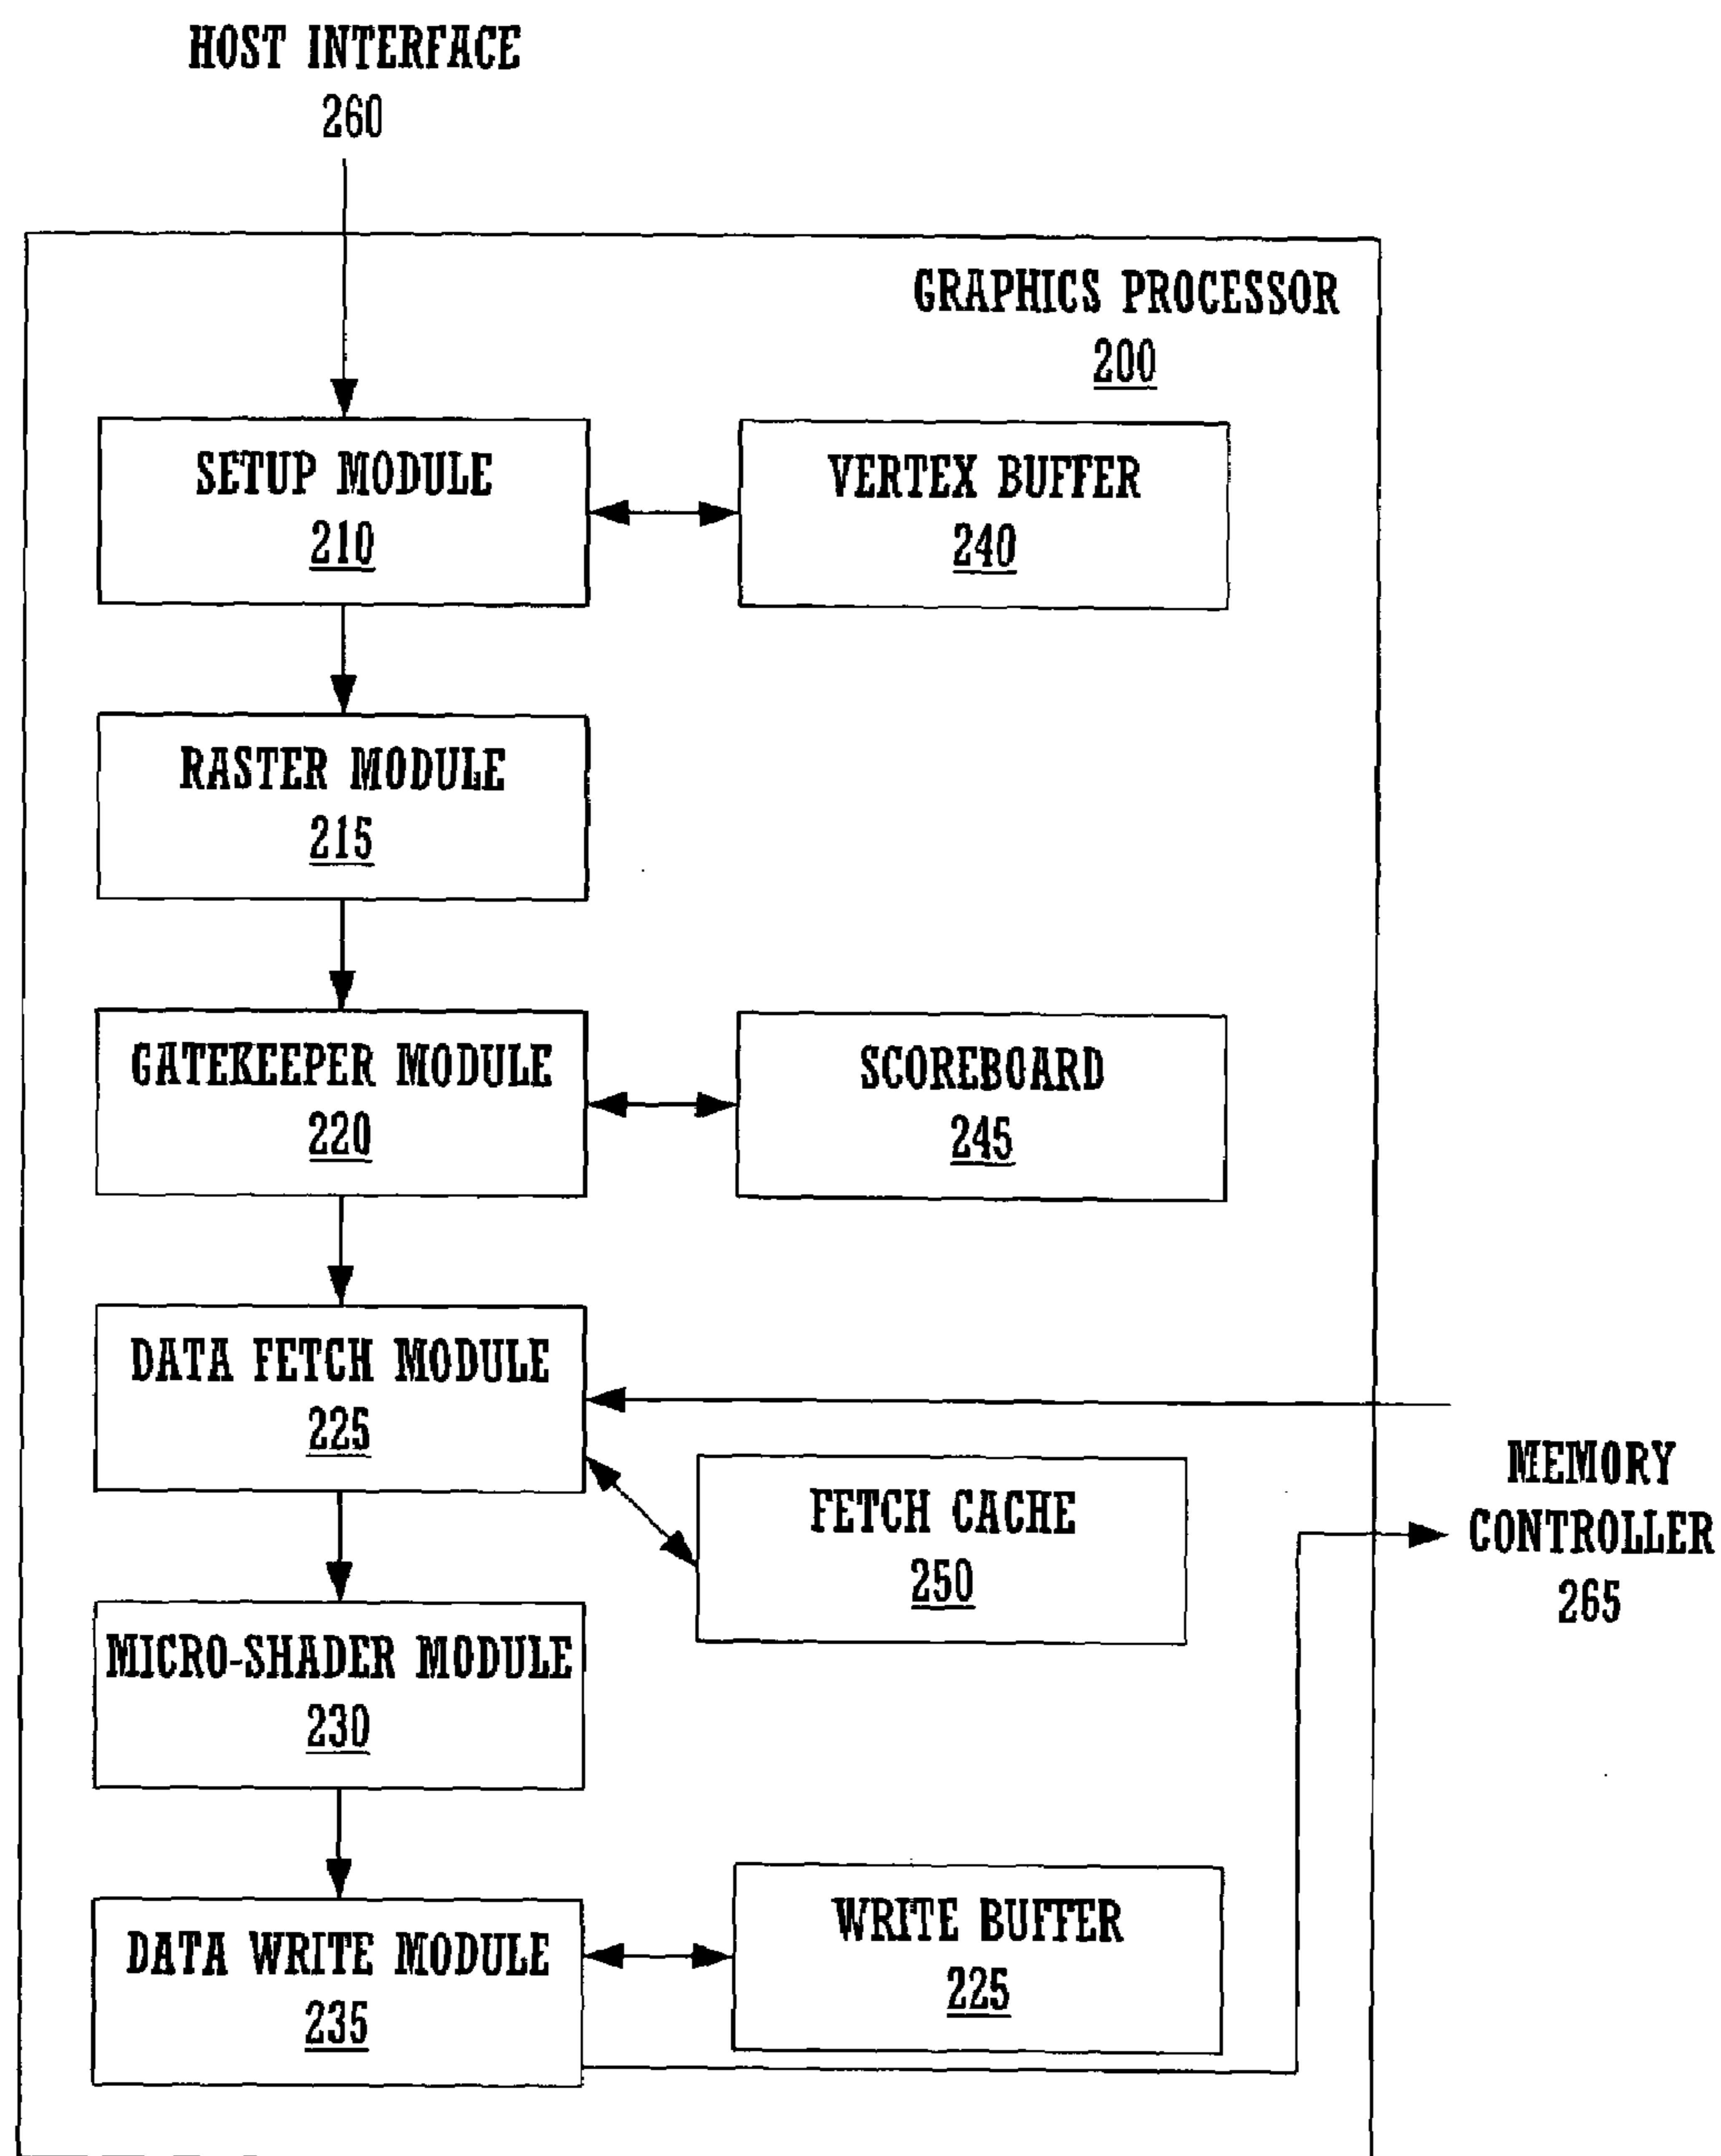


FIGURE 2

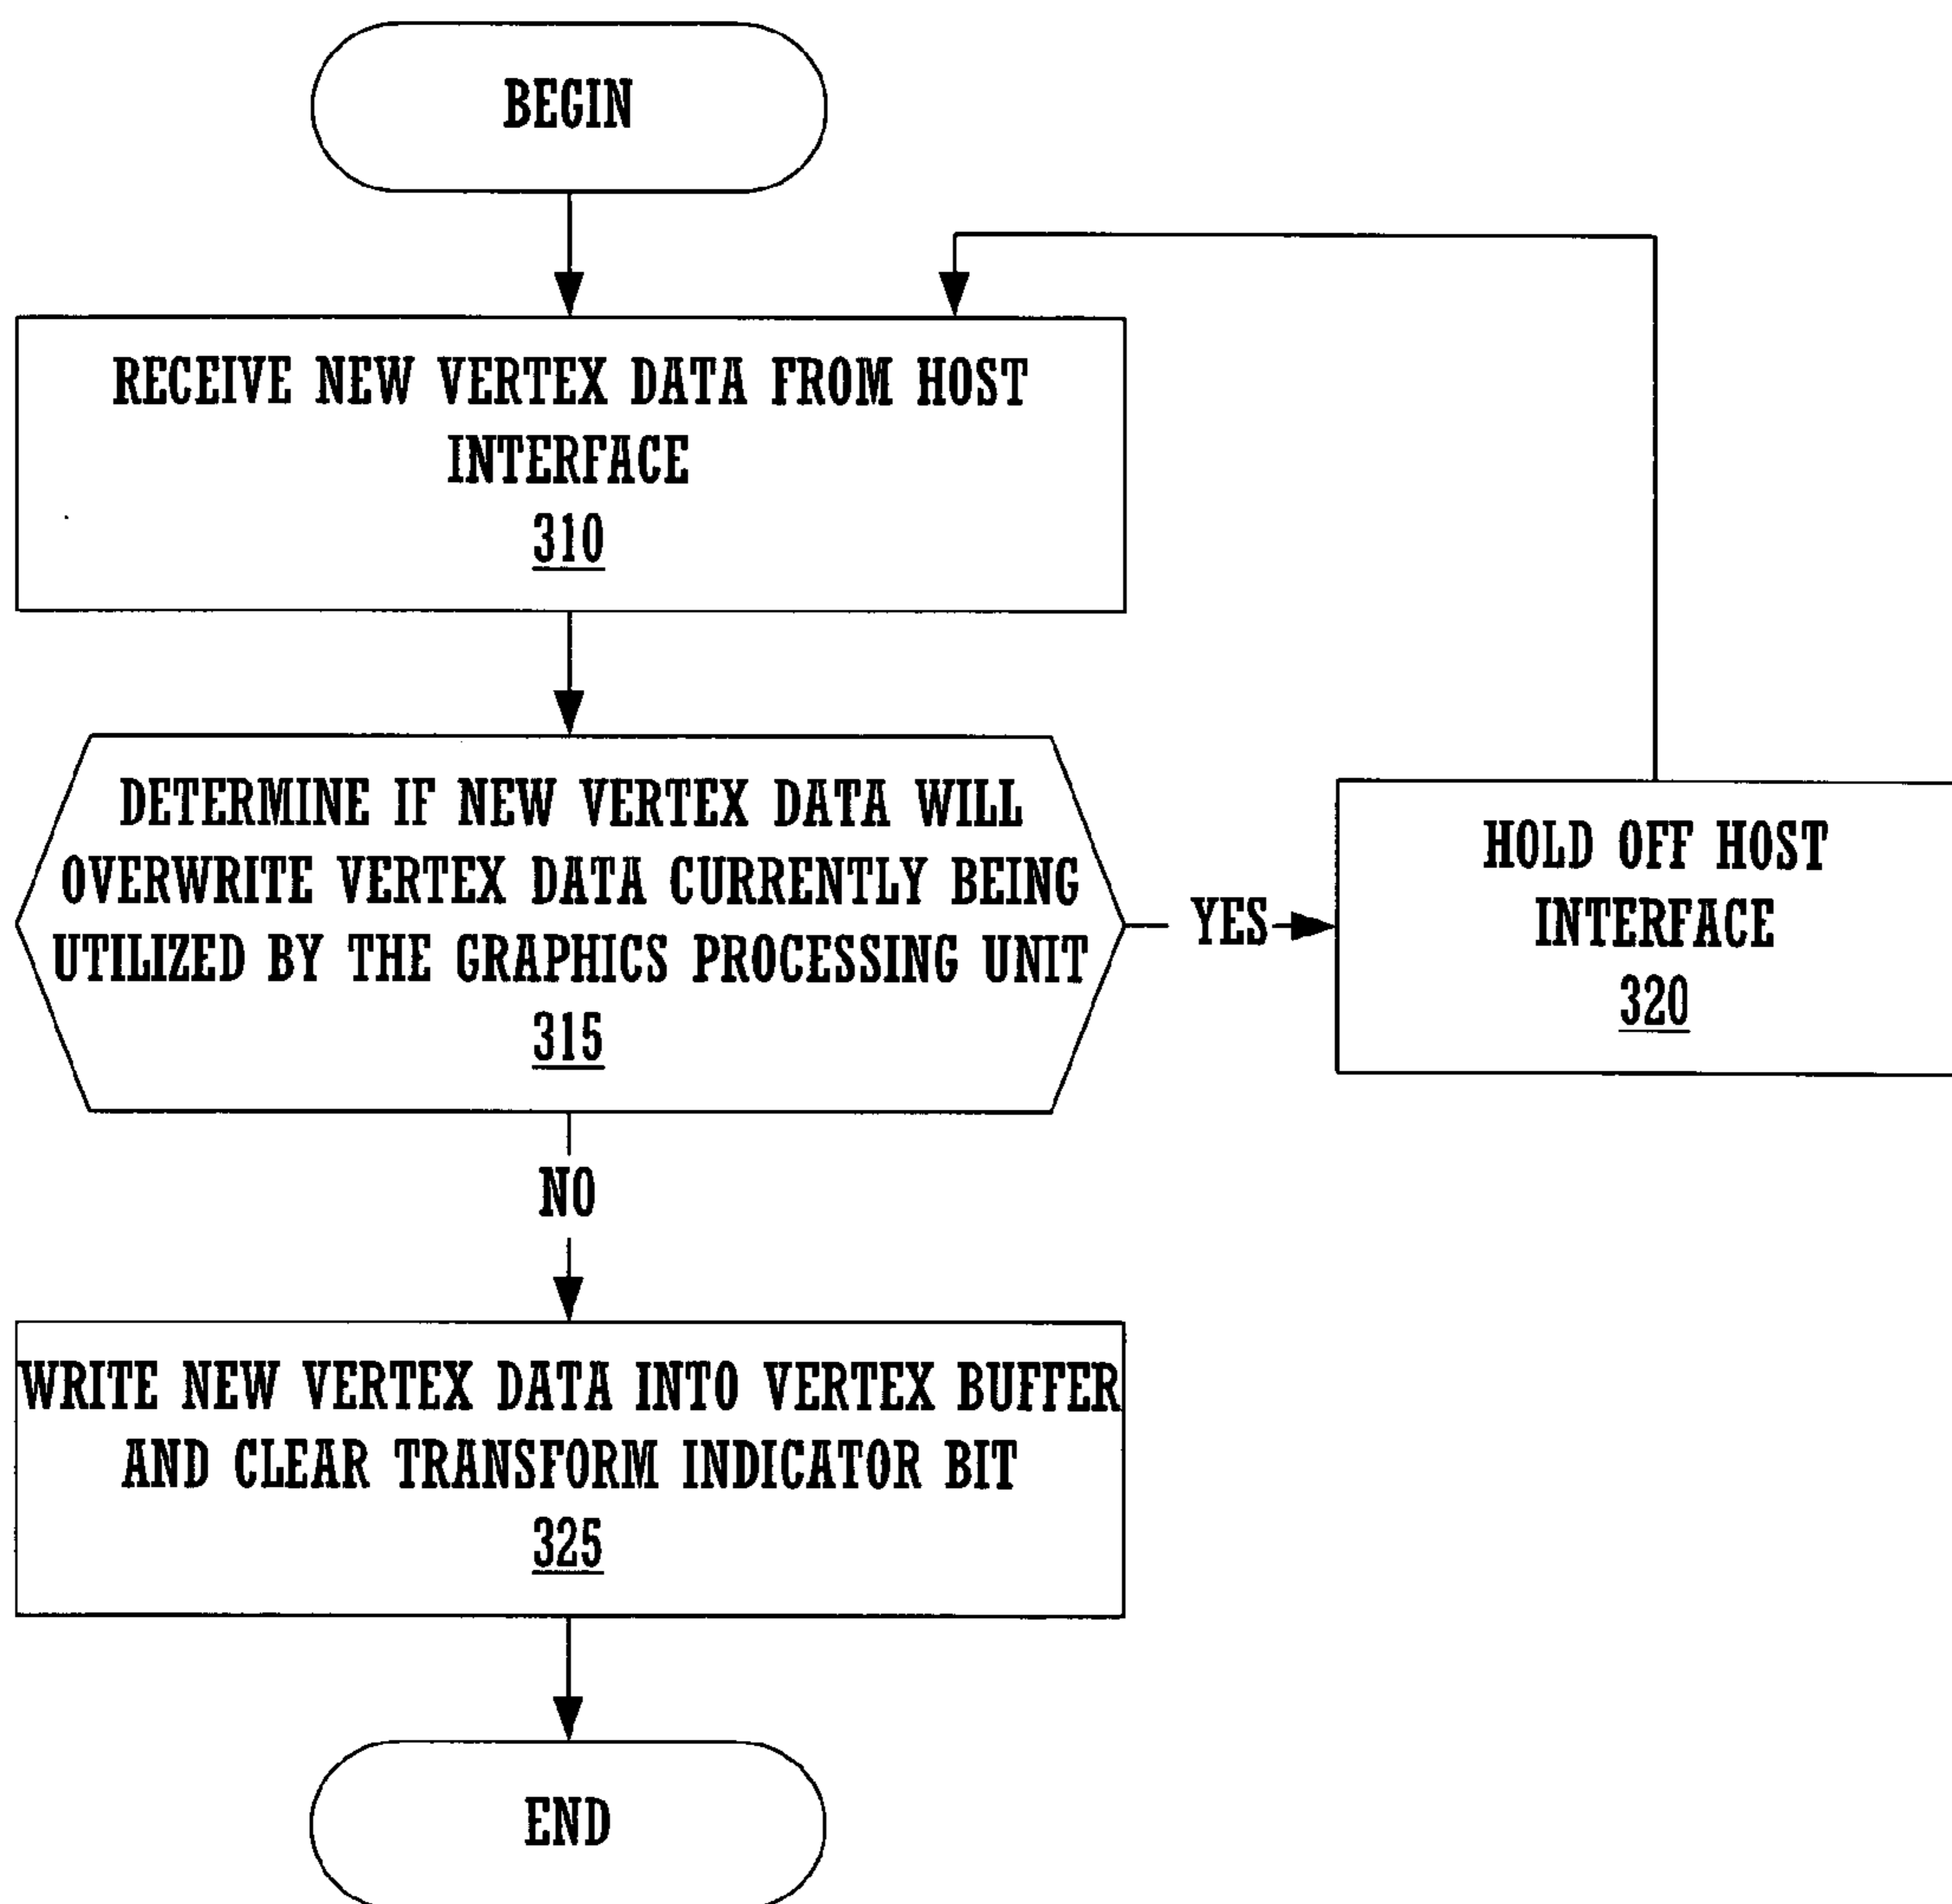


FIGURE 3A

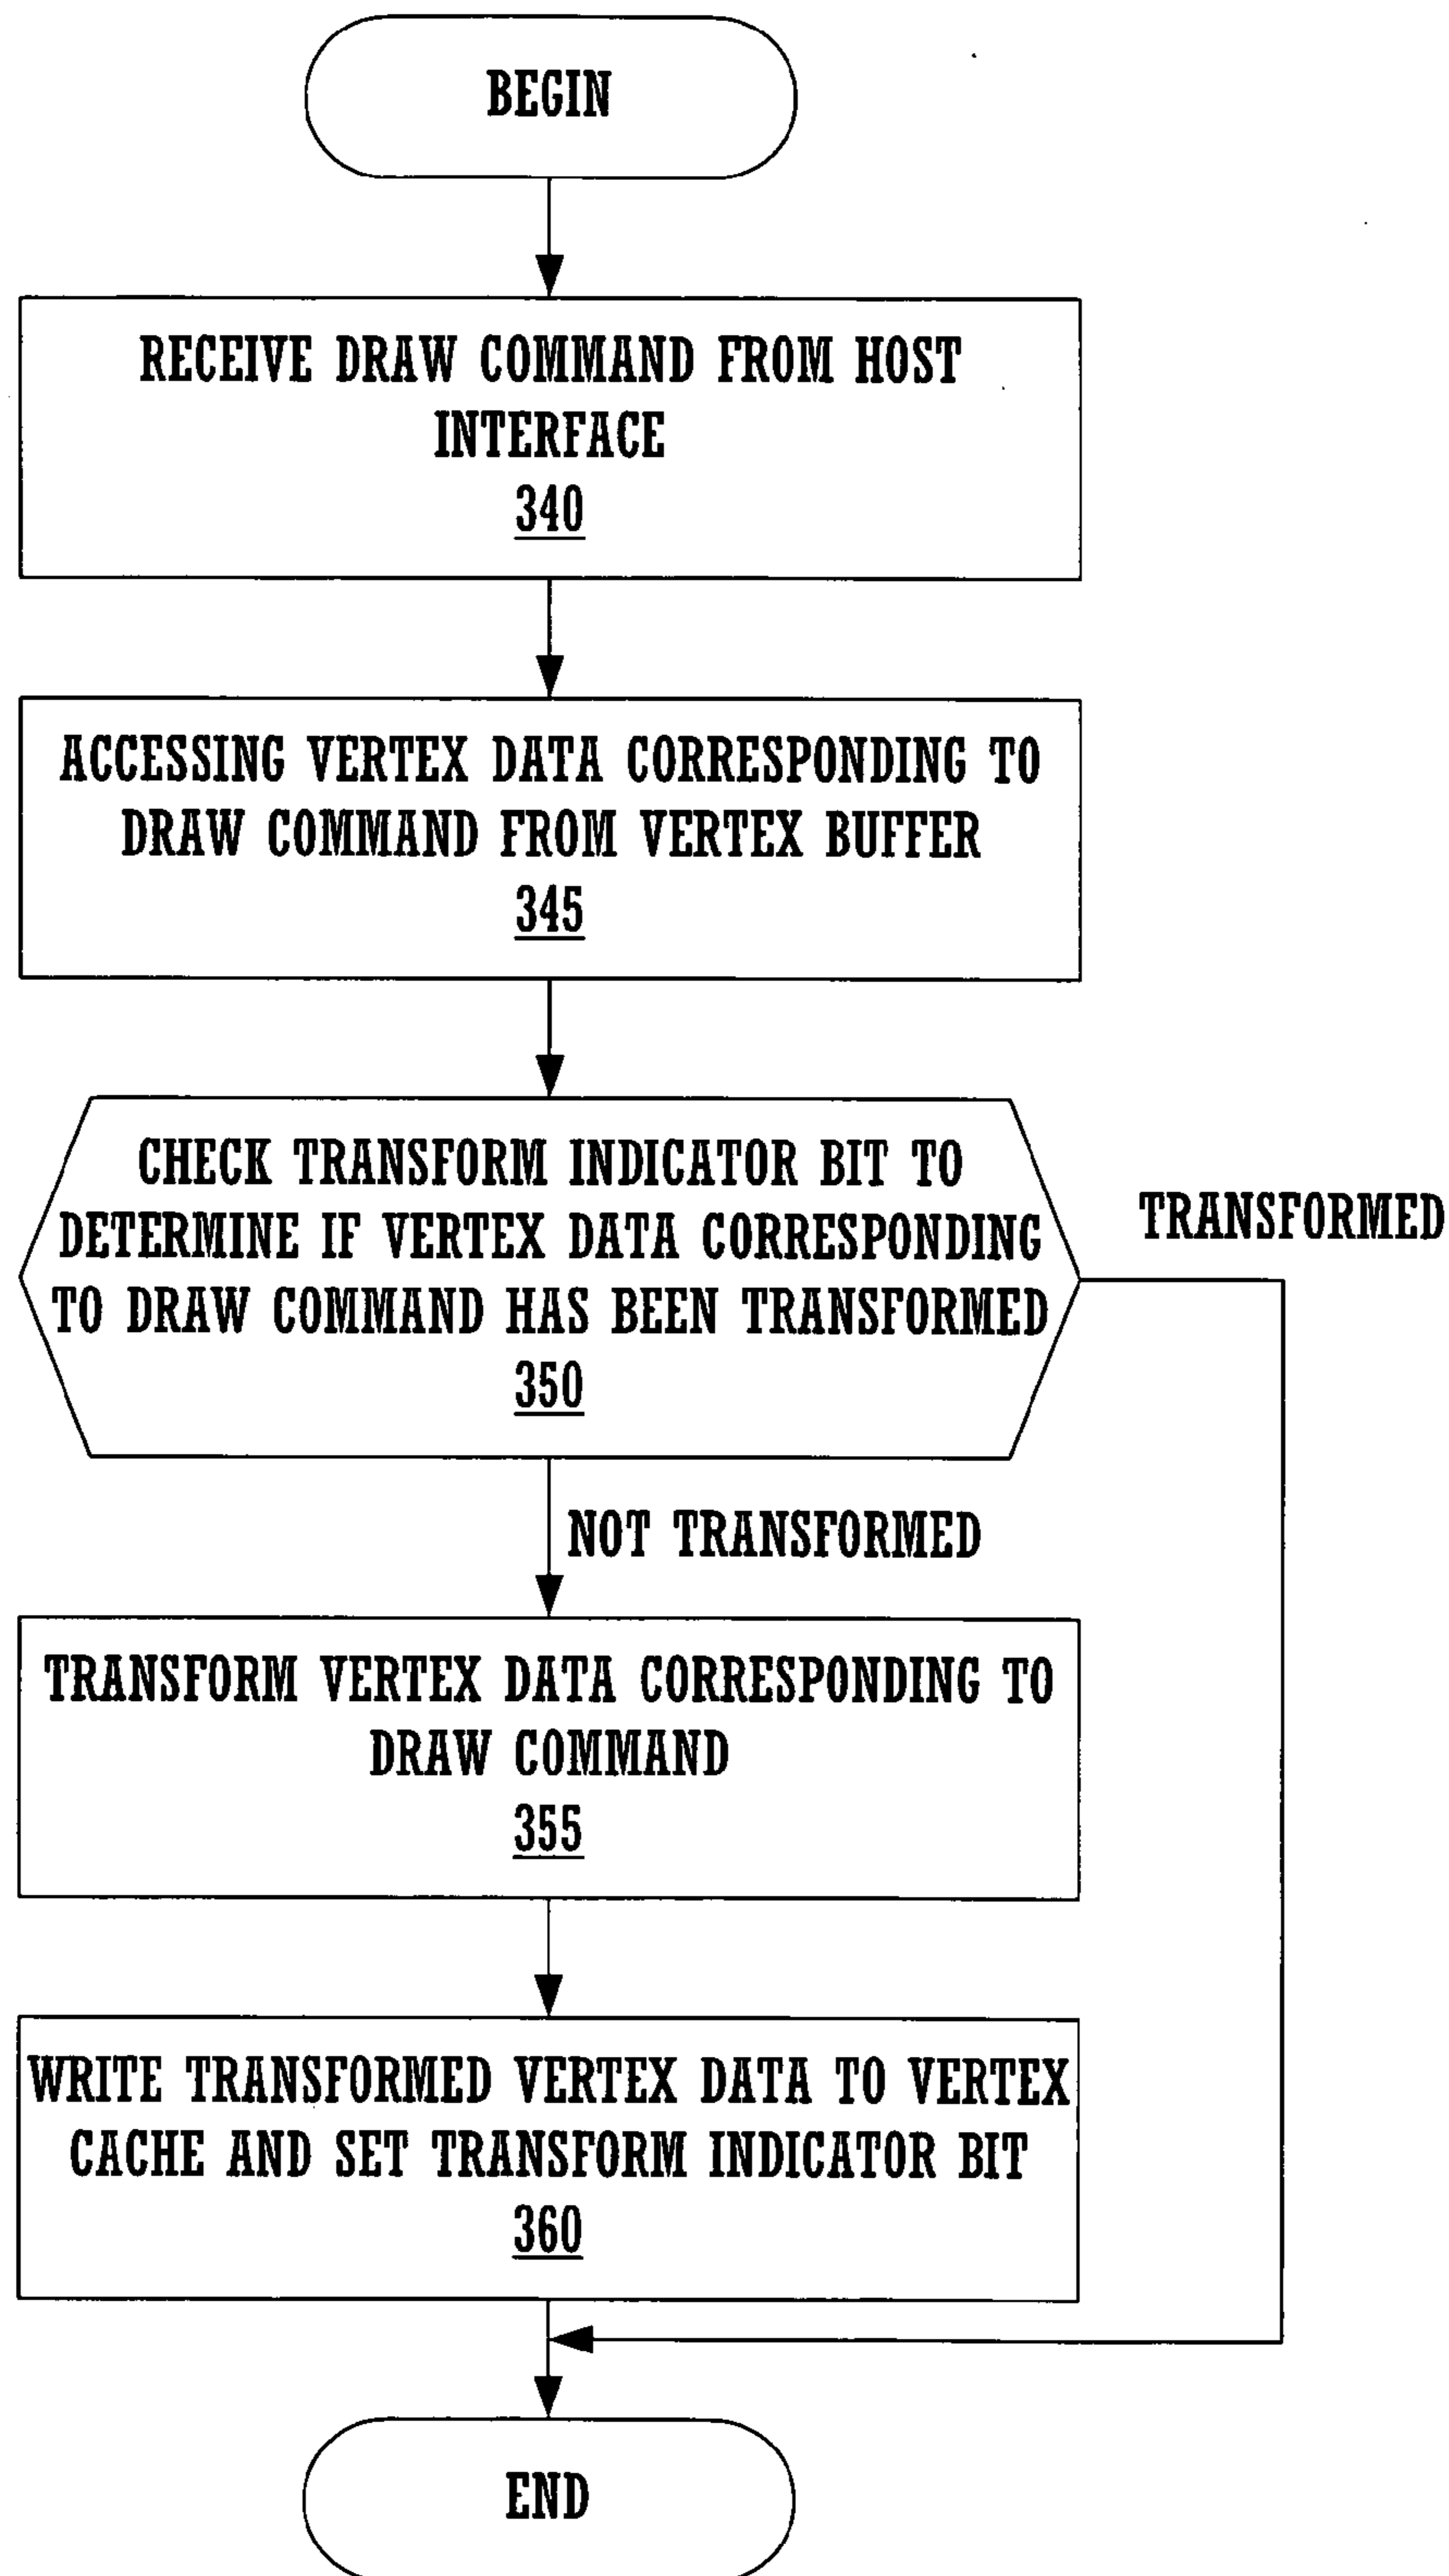


FIGURE 3B

400

405	1	VERTEX DATA 1	410
415	1	VERTEX DATA 2	420
425	1	VERTEX DATA 3	430
435	0	VERTEX DATA 4	440
445	0	VERTEX DATA 5	450
	.	.	
	.	.	
	.	.	
455	0	VERTEX DATA X	460

FIGURE 4A

400'

405	1	VERTEX DATA 1	410
415	1	VERTEX DATA 2	420
425	1	VERTEX DATA 3	430
435'	1	VERTEX DATA 4	440'
445	0	VERTEX DATA 5	450
	.	.	
	.	.	
	.	.	
455	0	VERTEX DATA X	460

FIGURE 4B

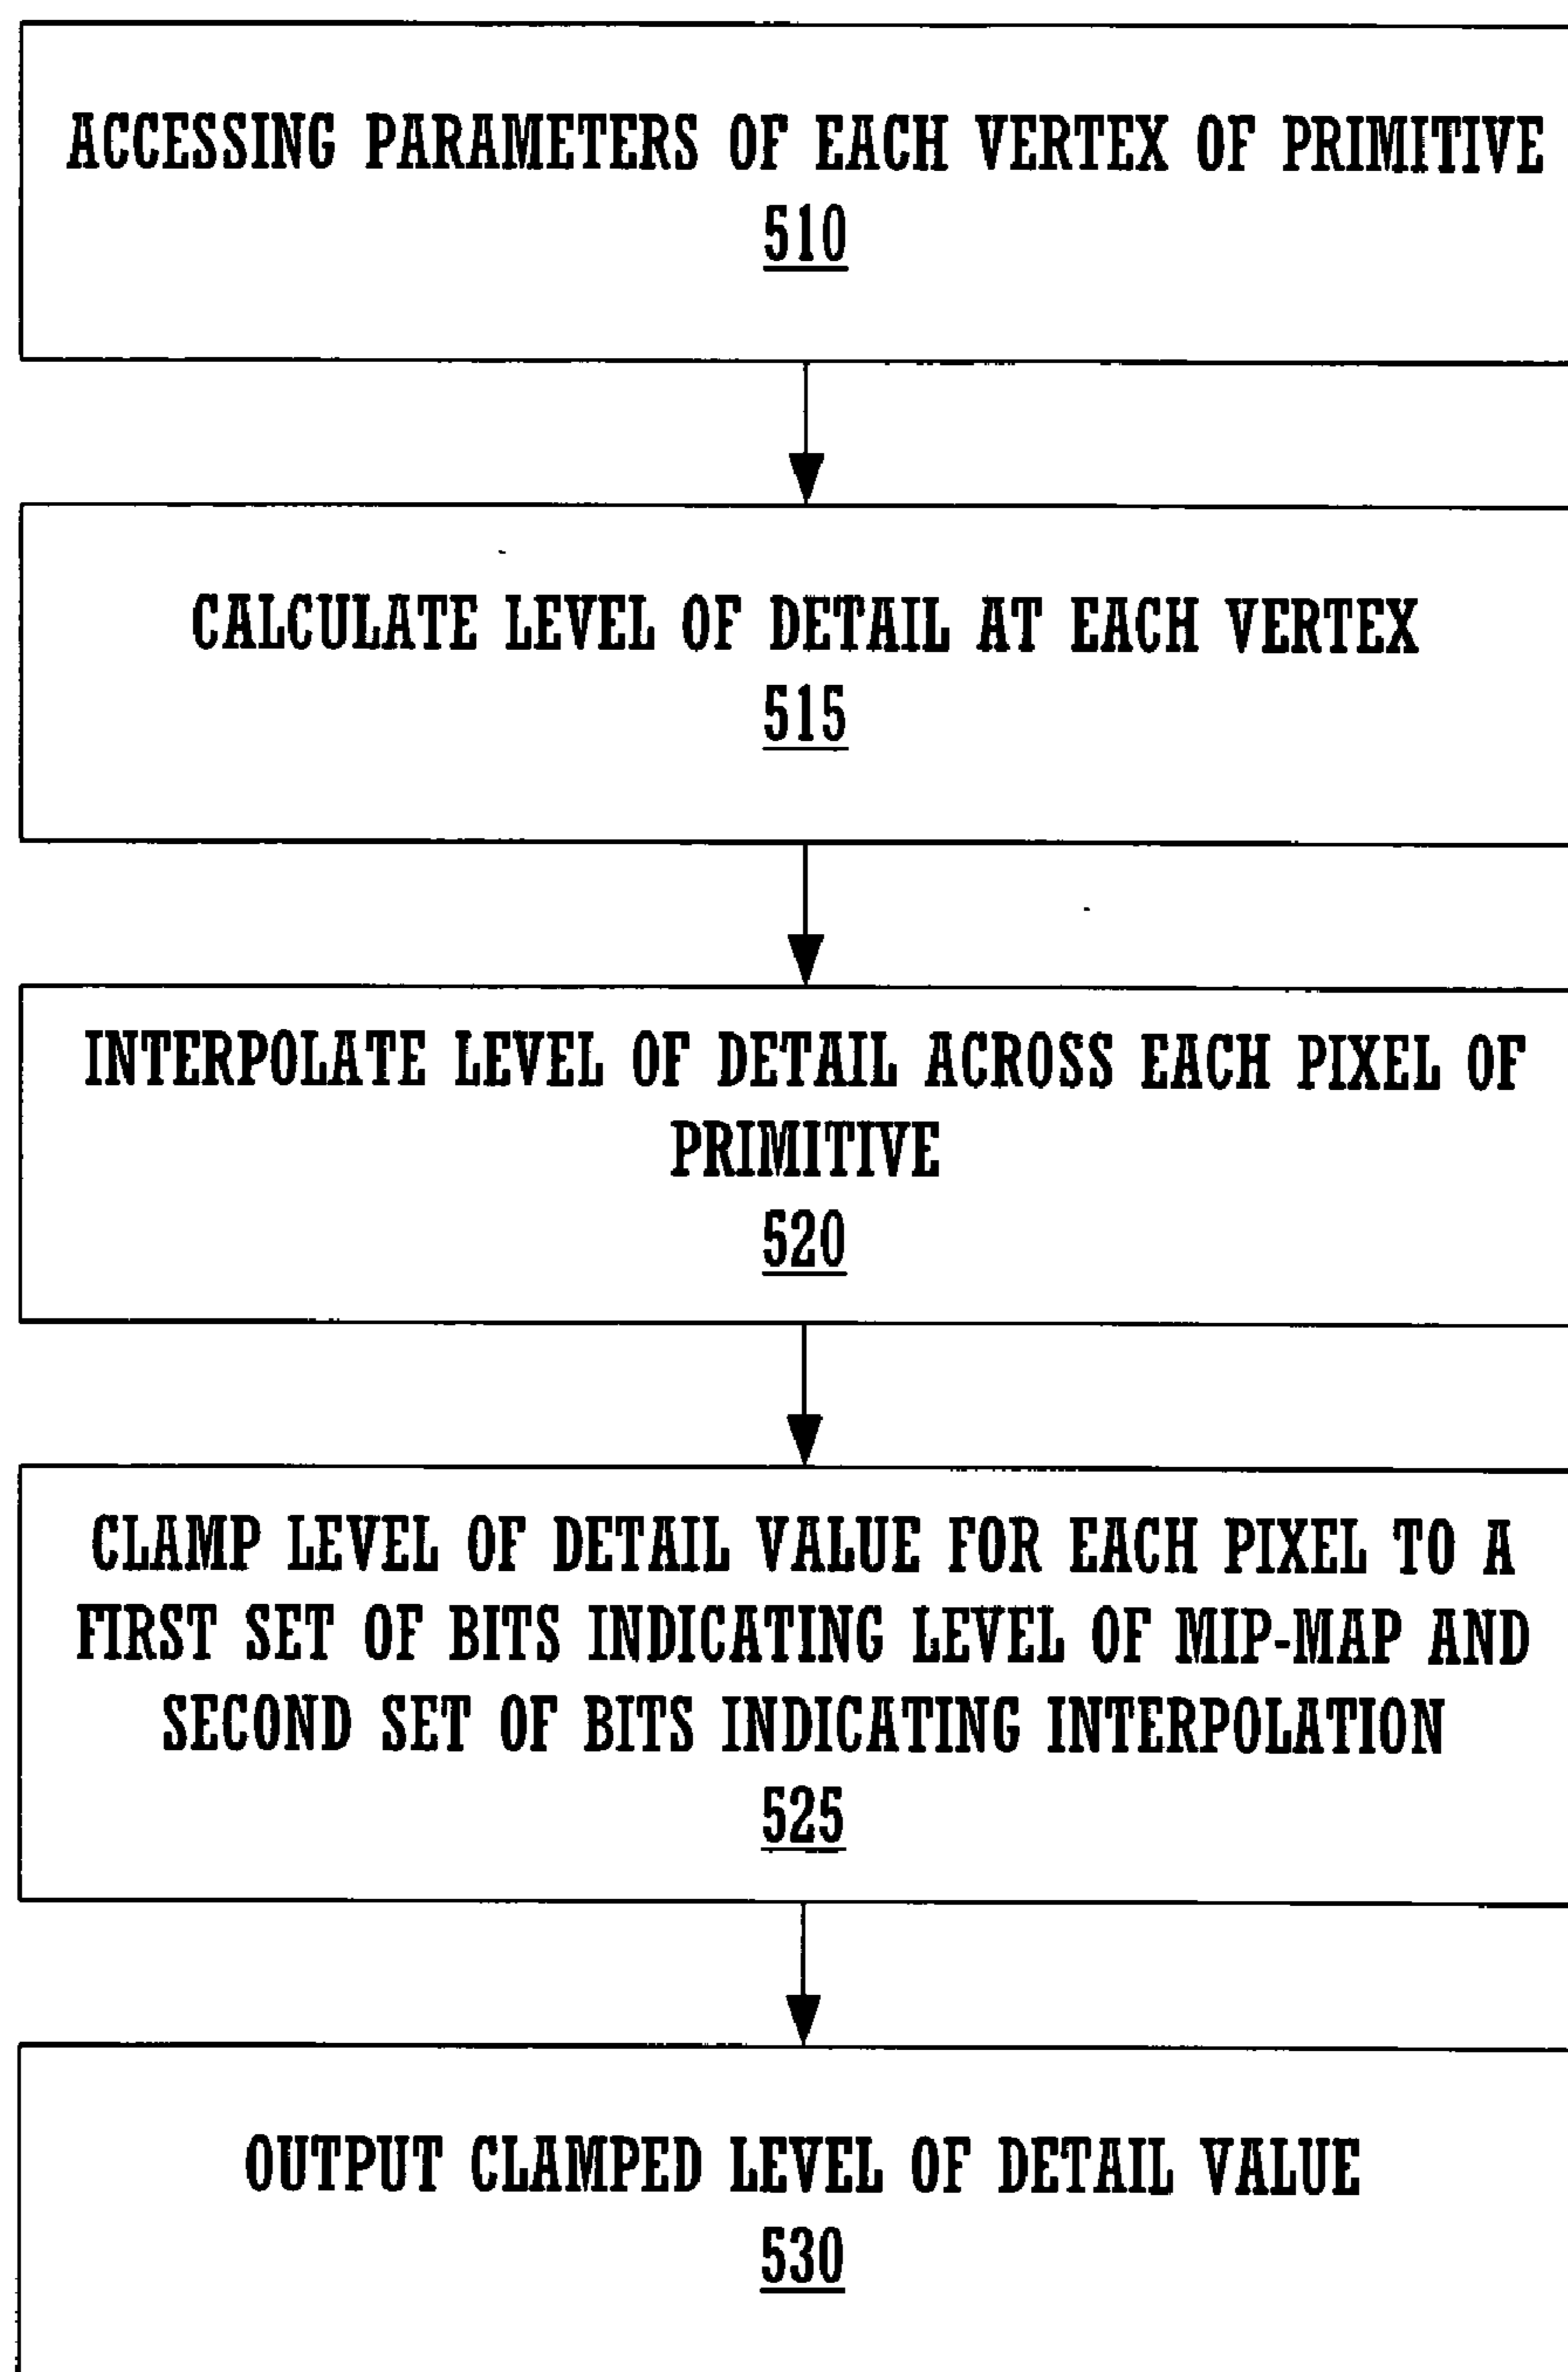


FIGURE 5

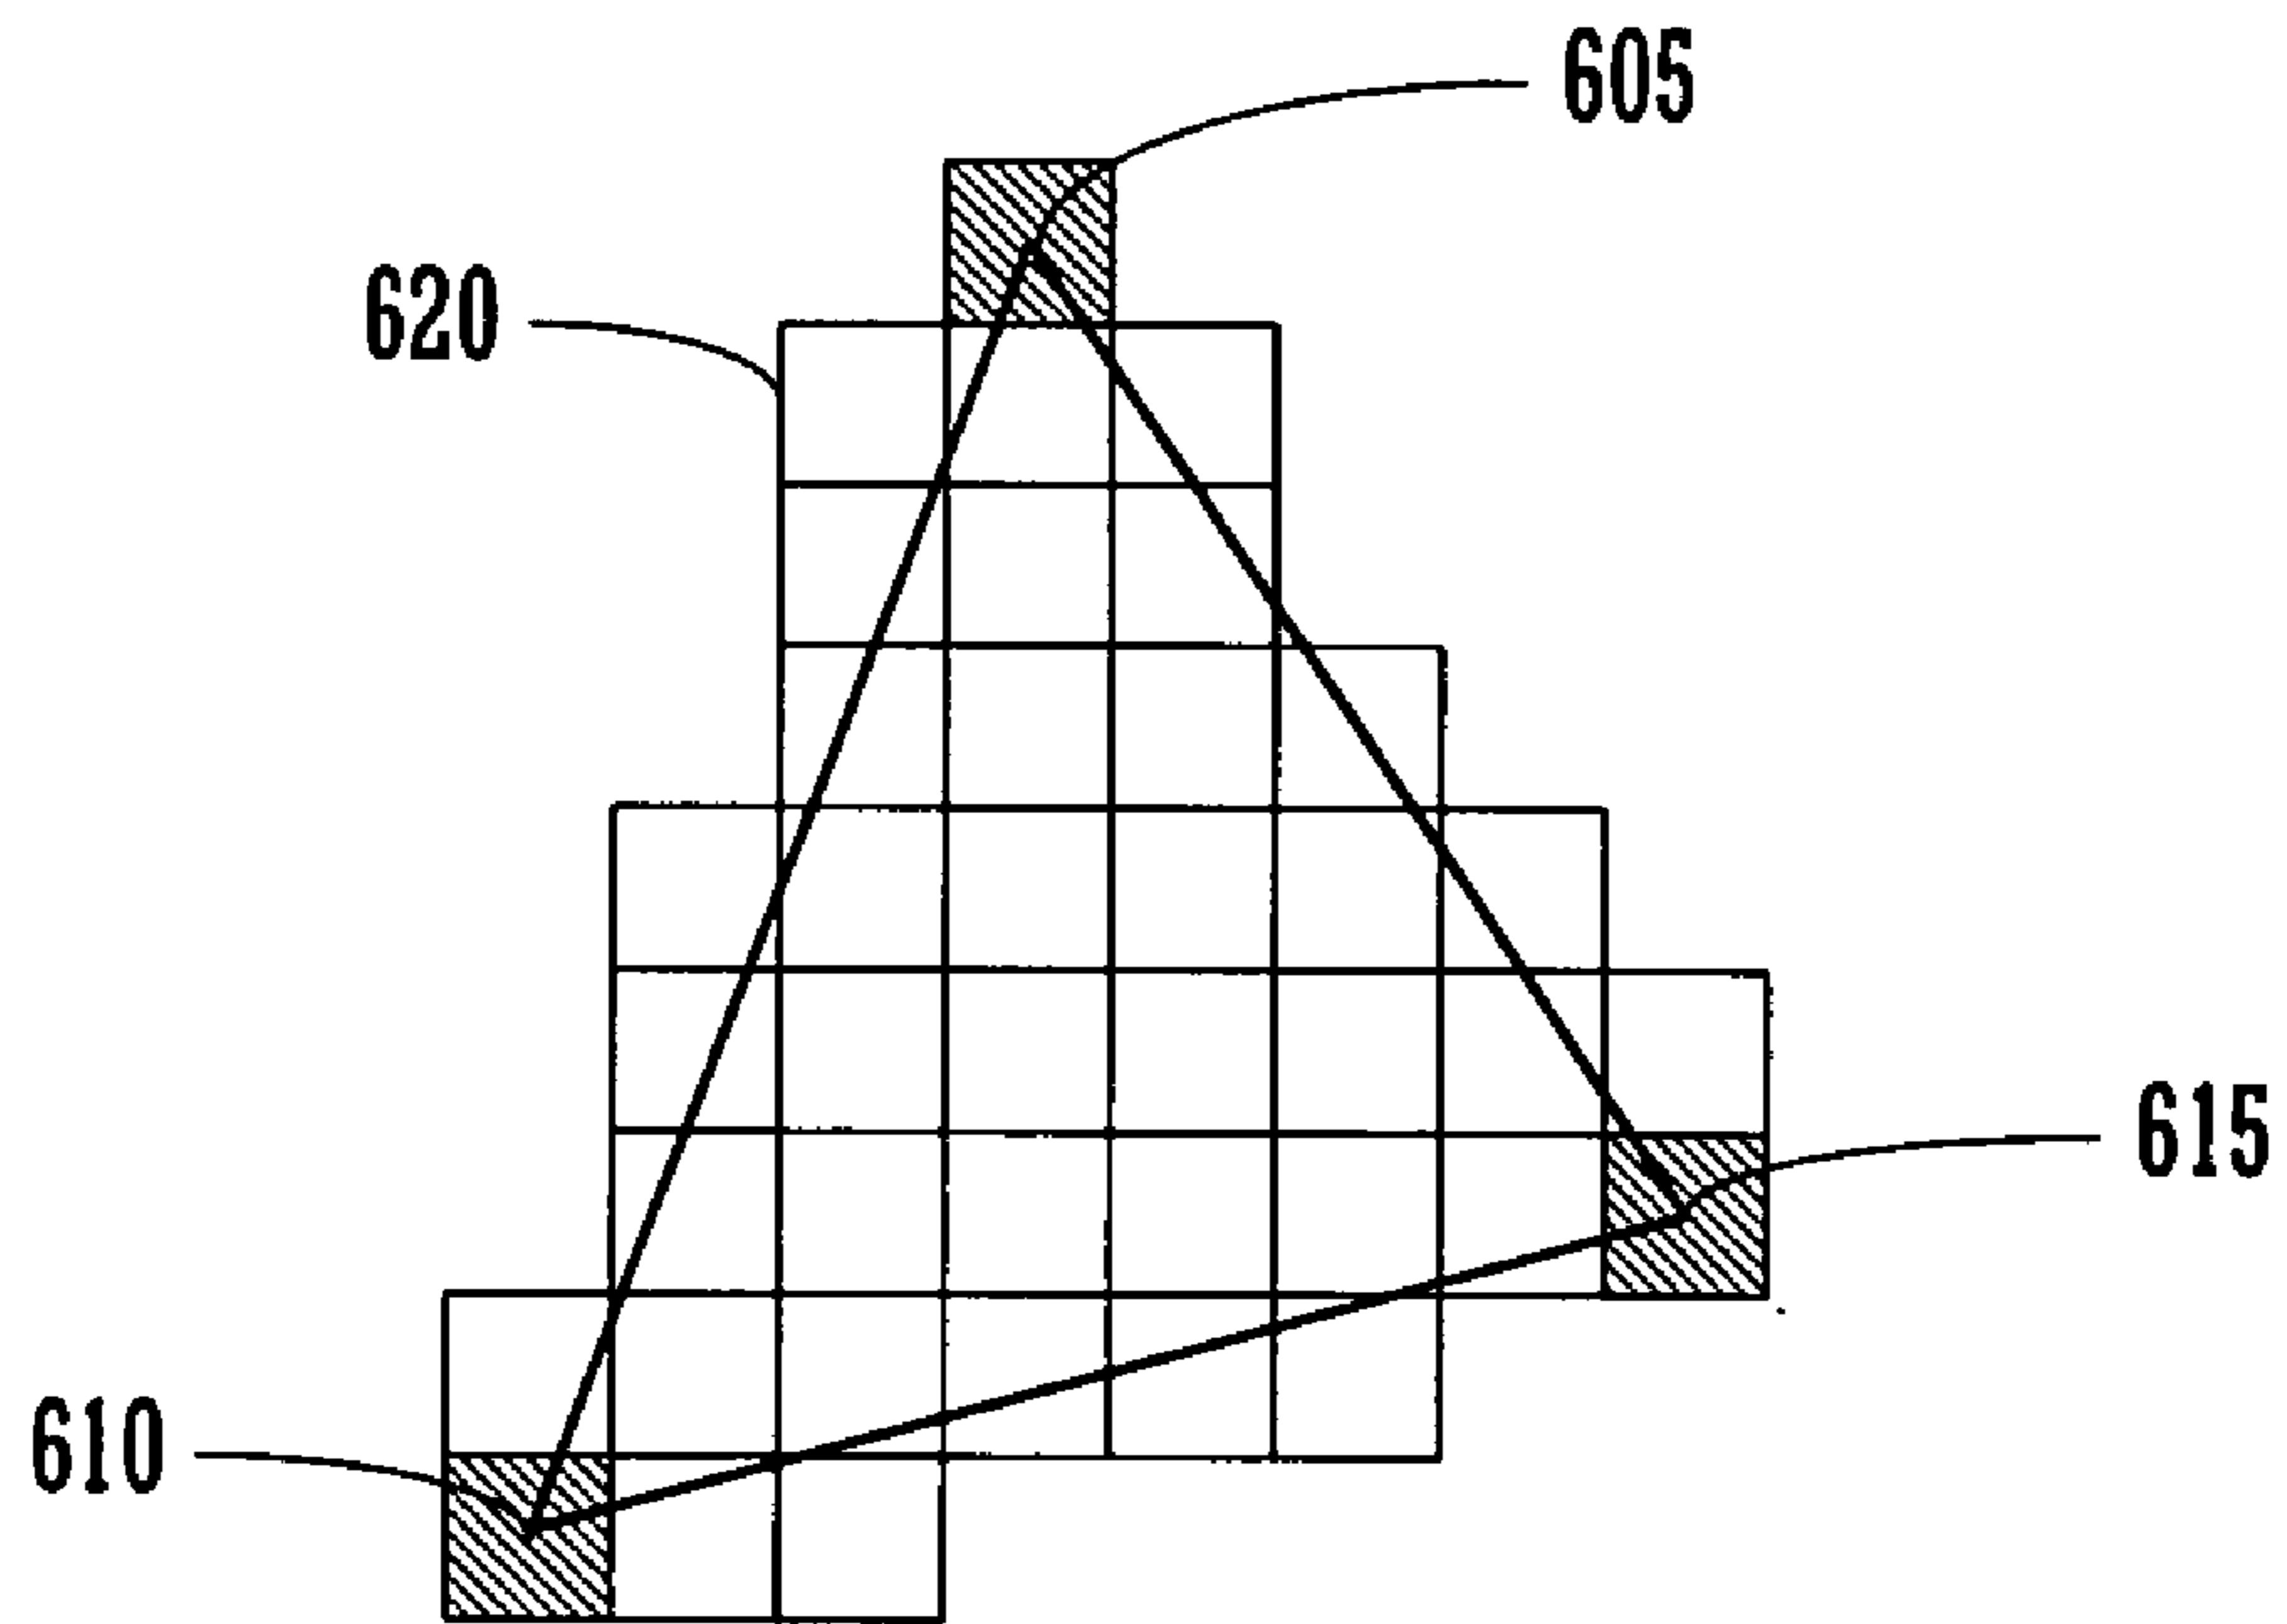


FIGURE 6

**METHOD AND SYSTEM FOR DEFERRED
EVALUATION OF TRANSFORMS IN
GRAPHICS PROCESSORS**

CROSS-REFERENCE TO RELATED
APPLICATIONS

The present application is related to U.S. patent application Ser. No. 10/846,786, filed May 14, 2004, by Edward Hutchins, entitled "Method and System for Interpolating Level-of-Detail in Graphics Processors," which is incorporated herein by reference.

BACKGROUND OF THE INVENTION

Three-dimensional graphics processing is utilized in a number of applications, from electronic games, and movies to computer aided design (CAD). Conventionally, three-dimensional graphics processing includes a multi-step rendering process of transitioning from a database representation of three-dimensional objects to a two-dimensional projection of the object into a display space. The process generally includes setting up a polygon model (e.g., a plurality of primitives) of objects, applying linear transformation to each primitive, culling back facing primitives, clipping the primitives against a view volume, rasterizing the primitives to a pixel coordinate set, shading/lighting the individual pixels using interpolated or incremental shading techniques, and the like. Typically, graphics processors are organized in a "deep" pipeline architecture, where each stage is dedicated to performing specific functions. A benefit of a deep pipeline architecture is that it permits fast, high quality rendering of even complex scenes.

The stages of a conventional graphics processor architecture are optimized for high-speed rendering operations (e.g., interpolating parameters, such as color, texture and depth over each two dimensional projection of a primitive). The architecture of the deep pipeline is configured in order to maximize the overall rendering throughput of the graphics processor. Generally, deep pipeline architectures have sufficient data throughput (e.g., pixel fill rate) to implement fast, high quality rendering on large display space devices of even complex scenes. For example, such conventional deep pipelines are configured to compute the various parameters required to render the pixels using multiple, high precision functions. The functions are implemented such that they generate high precision results even in those circumstances where such precision is redundant or unnecessary.

The dedicated stages of deep pipeline architected graphics processors require a relatively high transistor count. Accordingly, conventional graphic processors require a significant chip area, resulting in relatively high costs. In addition, to achieve fast, high quality rendering in deep pipeline architectures, various stages experience periods of idle processing cycles. Many of the stages consume about the same amount of power regardless of whether they are processing pixels or idle. Accordingly, conventional graphics processors consume significant power, even if the stages are performing comparatively little processing.

As a result of cost and power consumption considerations, conventional graphics processors are unsuitable for many mobile and wireless applications (e.g., wireless phones, personal digital assistants and the like). Therefore, what is desired is a graphics processor architecture having relatively low power consumption and costs.

SUMMARY OF THE INVENTION

Embodiments of the present invention are directed toward performing transforms on geometric primitives by a graphics processor. In one embodiment, a method of performing transformations on a geometric primitive includes storing a geometric primitive in a slot of a vertex buffer. The vertex buffer includes a plurality of slots, wherein vertex data associated with a particular geometric primitive may be stored. Each slot also stores a respective transform indicator bit. The method further includes clearing a transform indicator bit associated with said slot when the geometric primitive is stored. In response to a draw command involving a given geometric primitive, the given geometric primitive is transformed if the associated transform indicator bit is cleared. Otherwise, the given geometric primitive may be used without performing a transformation.

In another embodiment, a graphics processor includes a vertex buffer, a setup module and a raster module. The raster module is communicatively coupled to the setup module. The setup module is communicatively coupled to the vertex buffer. The vertex buffer stores a plurality of vertex data and a plurality of transform indicator bits, wherein each transform indicator bit is associated with a respective one of the vertex data. The setup module receives draw commands. The setup module transforms a given vertex data in response to the draw command when a corresponding transform indicator bit is in a first state. After performing the transformation, the transformed vertex data is saved in the vertex buffer and the corresponding indicator bit is set to a second state by the setup module. The raster module receives pixel packets that include the transformed vertex data and an instruction corresponding to said draw command from the setup module.

In yet another embodiment, a method of deferring evaluation of a transform includes buffering a plurality of vertex data. The method also includes receiving a draw command, accessing a given vertex data corresponding to the draw command and an associated transform indicator bit. The given vertex data is transformed if the associated indicator bit is cleared. After performing the transform, the vertex data is overwritten with the transformed vertex data and the associated transform indicator bit is set.

Embodiments of the present invention advantageously defer evaluation of transforms. Transforms are not automatically performed when vertex data is loaded into the vertex buffer. In addition, if a given vertex data has already been transformed, repeated transform of the given vertex data may be advantageously eliminated by embodiments of the present invention. Eliminating the repeated transform of vertex data currently stored in the vertex buffer reduces power consumed by the graphics processor. Furthermore, it is appreciated that the transform indicator bit enables storage of pre-transform vertex data and post-transform vertex data in a single vertex buffer. Storing vertex data (both pre- and post-transform) into a single vertex buffer simplifies arbitration between transform and setup units and allows the same underlying hardware to be utilized for both tasks. Hence, the hardware cost and complexity of the graphics processor may be reduced.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the present invention are illustrated by way of example and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

FIG. 1A shows a block diagram of exemplary computing device architecture for implementing embodiments of the present invention.

FIG. 1B shows a block diagram of another exemplary computing device architecture for implementing embodi- 5 ments of the present invention.

FIG. 2 shows a block diagram of a graphical processor, in accordance with one embodiment of the present invention.

FIGS. 3A and 3B show a flow diagram of an exemplary method of deferred transform evaluation, in accordance with 10 one embodiment of the present invention.

FIGS. 4A and 4B show an exemplary vertex buffer, in accordance with one embodiment of the present invention.

FIG. 5 shows a flow diagram of a method of interpolating level-of-detail (LOD), in accordance with one embodiment 15 of the present invention.

FIG. 6 shows an exemplary primitive, in accordance with one embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

Reference will now be made in detail to the embodiments of the invention, examples of which are illustrated in the accompanying drawings. While the invention will be 25 described in conjunction with these embodiments, it will be understood that they are not intended to limit the invention to these embodiments. On the contrary, the invention is intended to cover alternatives, modifications and equivalents, which may be included within the scope of the 30 invention as defined by the appended claims. Furthermore, in the following detailed description of the present invention, numerous specific details are set forth in order to provide a thorough understanding of the present invention. However, it is understood that the present invention may be 35 practiced without these specific details. In other instances, well-known methods, procedures, components, and circuits have not been described in detail as not to unnecessarily obscure aspects of the present invention.

Embodiments of the present invention provide a method 40 and system for implementing deferred evaluation of transforms in a graphics processor. Embodiments of the present invention enable the transfer of primitive parameters without automatically evaluating transforms thereon. Embodiments of the present invention also advantageously limit repetition 45 of evaluating transforms of the same parameters. The deferral of transform evaluation is achieved without a substantial increase in buffer capacity. Limiting the need to evaluate transforms reduces power expenditures and reduces the number of circuit elements. Embodiments of the present invention and their benefits are further described below.

Referring to FIGS. 1A and 1B, block diagrams of exemplary computing device architectures for implementing 55 embodiments of the present invention, are shown. The computing device may be a cellular telephone, PDA or other portable wireless appliance, navigation system (e.g., map enabled GPS), palm-sized computer, tablet computer, game console, personal entertainment center, media center PC, computer based simulator, desktop computer, laptop com- 60 puter, or the like. The computing device architectures provide the execution platform for implementing certain software-based functionality of embodiment of the present invention. As depicted in FIG. 1A, a first computing device may include a processor 105, system memory controller 110, system memory 115, host interface 120, graphics processor 125, graphics memory controller 130, graphics memory 135, display controller 140 and display 145. The processor 105

may be communicatively coupled to the system memory 115 through the system memory controller 110. The processor 105 may also be communicatively coupled to the graphics processor 125 through the host interface 120. The graphics processor 125 may be communicatively coupled to the graphics memory 135 through the graphics memory controller 130. The graphics memory controller 130 may also communicatively couple the display controller 140 to the graphics memory 135. The display 145 may be communi- 10 catively coupled to the display controller 140.

The processor 105 provides one or more applications by operating on instructions (e.g., computer executable code) and information (e.g., data) stored in the system memory 115 (e.g., computer-readable memory). The system memory controller 110 controls the flow of such information and instructions between the processor 105 and the system memory 115. Images to be output on the display 145 may be off-loaded to the graphics processor 125 by the processor 105. The images are off-loaded by transferring geometric 20 primitive parameters (e.g., vertex data representations of triangulated three-dimensional models of objects), draw commands and instructions for controlling the operation of the graphics processor 125. The primitive parameters, draw commands and instructions are transferred from the processor 105 to the graphics processor 125 under control of the host interface 120. In one implementation, a widely imple- 25 mented graphics application programming interface, such as the OpenGL-ES™ graphics language, Direct3D™, or the like, is utilized for the software interface between the graphics processor 125 and the applications provided by the processor 105.

The graphics processor 125 generates a color and depth value for each pixel of the display 145 in response to the received primitives, draw commands and instructions. In one implementation, the primitive parameters are stored in 35 a first buffer local to the graphics processor 125. The graphics processor 125 performs geometrical calculations (e.g., transforms) on the primitive in accordance with the draw commands. The graphics processor 125 also performs rasterization wherein values from the geometrical calculations are mapped to corresponding pixels in the display space. The graphics processor 125 also retrieves color, texture and depth information from the graphics memory 135. The graphics memory controller 110 controls the flow 45 of such information between the graphics memory 135 and the graphics processor 125. The color, texture and depth information may be cached local to the graphics processor 125.

The color and depth value of each pixel may then be 50 determined by the graphics processor 125 upon applying such functions as texture mapping, fogging, alpha testing, depth testing for culling occluded pixel, alphablend color combining and the like. The color of each pixel may then be cached in a second buffer local to the graphics processor 125. When the second buffer is filled, the output pixel data (e.g., color and depth value of each of a plurality of pixels) may then be stored by the graphics processor 125 in the graphics memory 135 under control of the graphics memory controller 130. The display controller 180 reads the output 60 pixel data out of the graphics memory 135, under control of the graphics memory controller 130. The display controller 180 generates scan control signals, from the pixel data, for output to the display 145. An image is project by the display 145 as a function of the scan control signals.

As depicted in FIG. 1B, a second computing device may include a processor 155, host interface 170, graphics processor 175, unified memory controller 160, unified memory

165, display controller 180 and display 185. The processor 155 may be communicatively coupled to the unified memory 165 through the unified memory controller 160. The processor 155 may also be communicatively coupled to the graphics processor 175 through the host interface 170. The graphic processor 175 may be communicatively coupled to the unified memory 165 through the unified memory controller 160. The unified memory controller 160 may also communicatively couple the display controller 180 to the unified memory 165. The display 185 may be communicatively coupled to the display controller 180.

The processor 155 provides one or more applications by operating on instructions (e.g., computer-executable code) and information (e.g., data) stored in the unified memory 165 (e.g., computer-readable memory). The unified memory controller 160 controls the flow of such information and instructions between the processor 155 and the unified memory 165. Images to be output on the display 185 may be off-loaded to the graphics processor 175 by the processor 155. The images are off-loaded by transferring geometric primitive parameters (e.g., vertex data), draw commands and instructions for controlling the operation of the graphics processor 175. The primitive parameters, draw commands and instructions are transferred from the processor 155 to the graphics processor 175 under control of the host interface 170. In one implementation, a widely implemented graphics application programming interface, such as the OpenGL-ES™ graphics library, Direct3D™, or the like, is utilized for the software interface between the graphics processor 175 and the applications provided by the processor 155.

The graphics processor 175 generates a color and depth value for each pixel of the display 185 in response to the received primitives, draw commands and instructions. In one implementation, the primitive parameters are stored in a first buffer local to the graphics processor 175. The graphics processor 175 performs geometrical calculation (e.g., transforms) on the primitive in accordance with the draw commands. The graphics processor 175 also performs rasterization wherein values from the geometrical calculations are mapped to corresponding pixels in the display space. The graphics processor 175 also retrieves color, texture and depth information from the unified memory 165. The unified memory controller also controls the flow of such information between the graphics processor 175 and the unified memory 165. The color, texture and depth information may be cached local to the graphics processor 175.

The color and depth values of each pixel may then be determined by the graphics processor 175 upon applying such functions as texture mapping, fogging, alpha testing, depth testing for culling occluded pixels, alphablend color combining and the like. The color of each pixel may then be cached in a second buffer local to the graphics processor 175. When the second buffer is filled, the output pixel data (e.g., color and depth value of each of a plurality of pixels) may then be stored by the graphics processor 175 in the unified memory 165, under control of the unified memory controller 160. The display controller 180 reads the output pixel data output from the unified memory 165, under control of the unified memory controller 160. The display controller 180 generates scan control signals, from the pixel data, for output to the display 185. An image is projected by the display 185 as a function of the scan control signals.

Referring now to FIG. 2, a block diagram of a graphical processor 200, in accordance with one embodiment of the present invention, is shown. As depicted in FIG. 2, the graphical processor 200 includes a setup module 210, a raster module 215, a gatekeeper module 220, a data fetch

module 225, a micro-shader module 230, and a data write module 235. The graphics processor 200 may further include a vertex buffer 240, a scoreboard 245, a fetch cache 250 and/or a write buffer 255. The modules 210–235 of the graphical processor 200 are arranged in a pipeline architecture.

The setup module 210 may be communicatively coupled to a host interface 260. The setup module 210 receives geometric primitives (e.g., parameters defining a polygon in three-dimensional space), draw commands and instructions (e.g., for controlling operation of the modules 210–253) from the host interface 260. The setup module 210 prepares parameters (e.g., vertex data) of the primitives for rasterization. In one implementation, the primitives include parameters describing the three vertices of a triangle in three-dimensional space. The setup module 210 supports input formats of 32 bit IEEE floating point, 32 bit S15.16 fixed point (e.g., OpenGL-ES™ fixed standard) and/or 8 bit packet (e.g., for colors) data packets. The data packets may contain up to 24 parameters in addition to x, y, z, and w parameters, per primitive.

The setup module 210 may be communicatively coupled to the vertex buffer 240. The vertex buffer 240, in one implementation, includes storage (e.g., a plurality of discrete slots) for 256 words (e.g., vertex words). In one implementation, the discrete slots may either be in 8 or 16 word vertex formats, under control of software. Primitive parameters (e.g., three-dimensional vertex data) may be loaded from the host interface 260 into the vertex buffer 240 through the setup module 210, under software control. Although the host interface 260 is communicatively coupled to the setup module 210, the host interface 260 may effectively load vertex data directly into the vertex buffer 240.

A transform (e.g., matrix multiplication) is performed upon the vertices to convert the three-dimensional arbitrary space that a given application is working in to the two-dimension space that is being projected onto (e.g., view surface). The transform process is not automatically performed when vertex data is loaded into the vertex buffer 240 by the host interface 260.

Referring now to FIGS. 4A and 4B, an exemplary vertex buffer 240, in accordance with one embodiment of the present invention, is illustrated. As depicted in FIG. 2 in conjunction with FIGS. 4 and 4B, a transform indicator bit 405, 415, 425, 435, 445 and 455 is associated with each vertex word 410, 420, 430, 440, 450 and 460, respectively, stored in the vertex buffer 240. The associated transform indicator bit 405 indicates whether the given vertex word 410 has been transformed. When the host interface 260 loads a vertex word (e.g., 440) into the vertex buffer 240, the data is in three-dimensional space and therefore the associated transform indicator bit (e.g., 435) is set to a first state (e.g., cleared). When the setup module 210 accesses a draw command from the host interface 260, the setup module 210 retrieves vertex data, applicable to the given draw command, from the vertex buffer 240. The state of the associated transform indicator bit is checked upon retrieving the given vertex data. If the given vertex word is three-dimensional data (e.g., 440) as indicated by the transform indicator bit, the setup module 210 transforms the given vertex data from a three-dimensional into a two-dimensional representation. The vertex word (e.g., 440) may be over written with the transformed (e.g., two-dimensional) vertex word (e.g., 440') and the associated transform indicator bit (e. 435') is set to a second state (e.g., set). For a given triangle, if its indicator bit indicates transformed data, then the triangle is accessed directly for rasterization without transformation required.

After transformation, the setup module stores the vertex data back to the vertex buffer and sets the indicator bit to indicate the transformation state.

Referring now to FIG. 6, an exemplary primitive, in accordance with one embodiment of the present invention, is illustrated. As depicted in FIG. 2 in conjunction with FIG. 6, the level-of-detail (LOD) for each vertex 605–610 of the given primitive is also calculated by the setup module 210. Calculation of the LOD includes determining the approximate area of each vertex pixel projected into the image space (e.g., the area of each box represents an area of an exemplary pixel projected into the image space). In the conventional art, an LOD value of each vertex of a primitive may be calculated as a function of the primitive's vertex X coordinate value, the vertex Y coordinate value, the texture S coordinate value and the texture T coordinate value according to Equation 1, as follows:

$$LOD = \text{Log}_2(\max(\sqrt{ds/dx^2 + dt/dx^2}, \sqrt{ds/dy^2 + dt/dy^2})) \quad (1)$$

In accordance with embodiments of the present invention, LOD interpolant values (LOD_interp(n)) of the vertices of a triangular primitive are determined prior to calculating the logarithm base two. The pre-Log₂() LOD is also divided by the square of the homogeneous view distance (w) square for each vertex (n) according to Equation 2, as follows:

$$LOD_interp(n) = [\max(\sqrt{ds/dx(n)^2 + dt/dx(n)^2}, \sqrt{ds/dy(n)^2 + dt/dy(n)^2})] / w(n)^2 \quad (2)$$

It is appreciated that the pre-Log₂() LOD should actually be divided by w(n)⁴ (e.g., based on the partial derivatives of the pre-Log₂() LOD calculations with respect to changes in X and Y). However, using w(n)² is sufficiently accurate to provide relatively good quality images as perceived by the human eye under substantially all circumstances. Utilizing w(n) advantageously reduces computation costs and the dynamic range. Upon determining the LOD interpolant values of the vertices, they may be interpolated just like any other perspective-correct primitive parameter.

Accordingly, the LOD calculations are done on a per primitive basis, as opposed to per-pixel. Calculating the LOD on a per-primitive basis reduces power consumption by the graphics processor 200. The LOD values of the vertex pixels may then be utilized in the raster module 215 to interpolate the LOD values of the other pixels in the given primitive based on the vertex data.

In one implementation, the setup module 210 calculates barycentric coefficients for rastering. In a barycentric coordinate system, distances in primitives (e.g., triangles) are measured with respect to its vertices. The use of barycentric coordinates reduces the required dynamic range, which permits using fixed-point calculations that require less power than floating point calculations. The two-dimensional parameters (e.g., barycentric) are sent from the setup module to the raster module. The setup module 210 may also perform clipping and viewport operations. The setup module 210 may also perform culling of back-facing primitives and occlusion. Thus, culling and occlusion operations are performed before texturing and shading. Eliminating pixels, as a result of culling and/or occlusion, before texturing and shading reduces power consumption and reduces the associated memory accesses by the data fetch module 225 and the data write module 235. The setup module 210 may also support guard-band clipping, power management for the rest of the pipeline, and/or collection of statistics from other modules 215–235, buffers 240, 255, caches 250 and/or the scoreboard 245.

The raster module 215 may be communicatively coupled to the setup module 210. The raster module 215 receives instructions and applicable transformed vertex data and performs rasterization thereupon. The raster module 215 determines which pixels on a screen correspond to a given polygon as defined by the transformed vertex data. Rasterizing includes “walking” the primitive and generating pixel packets that contain parameters (e.g., set of descriptions for each pixel). The raster module 215 also determines shader processing operations to be performed, as part of rendering, on each pixel, such as color, texture and fog operations. Hence, the raster module 215 converts from per polygon (e.g., triangle) to per pixel processing.

The raster module 215 generates at least one pixel packet for each pixel of a primitive. Each pixel packet includes fields for a payload of pixel parameters required for processing, an instruction sequence of operation to be performed on the pixel packet, and sideband information. The payload of the pixel may include color values, texture coordinates (e.g., S and T), X coordinate values, Y coordinate values, Z depth values, LOD values and the like. In one implementation, the color field may be a pair of 8-bit precision color values, or packed 5555 ARGB color values. The texture field may be 16-bit S and T coordinates, with a 4-bit LOD value. The depth field may be a 16-bit value. In one implementation, the sideband information includes a valid field, a kill field, a packet type descriptor (whether the packet is a register write or contains pixel data) or the like. The generated pixel packets may be four 20-bit variables per row, and up to 4 rows. The generated pixel packets may also include 4 high-precision and 4 low-precision perspective-corrected iterated values per row.

The raster module 215 interpolates the LOD of the pixels from the calculated LOD interpolant values of the primitive vertices 605–615. However, a straight interpolation across the primitive (e.g., triangle) will provide relatively poor image quality. Accordingly, the final LOD (final_LOD(i)) for each pixel is interpolated according to Equation 3:

$$\text{final_LOD}(i) = \text{Log}_2(\text{LOD_interp}(n)) * W^2 \quad (3)$$

where final_LOD(i) are interpolated values of LOD for a respective pixel, i, of the pixels. Equation 3 may be evaluated for each texture to be applied to the pixels of the primitive. Assuming that a fast approximation of the logarithm base two is available, Equation 3 can be simplified as given by Equation 4:

$$\text{final_LOD}(i) = \text{approx_Log}_2(\text{LOD_interp}) + \text{Log}_2(w * 2) \quad (4)$$

The final LOD of each pixel is clamped between zero and Log₂(N), where N is the maximum dimension of the largest texture supported by the device. The LOD value has an integer portion and a fractional portion. In one implementation, an eight-bit value representing the LOD of each pixel, where four integer bits indicate which of a plurality of mip-map levels to use and four fractional bits are utilized to interpolate the mip-map level.

It is appreciated that if there are many texels (texture elements) in a given pixel, than an average texel (e.g., filtered-value) may be determined. In mip-mapping the texture is pre-filtered into a plurality of resolutions. For example, a first file of the texture has a first resolution, a second file has a quarter of the resolution of the primary texture, a third file has a sixteenth of the resolution, and so on until a last file contains a single value (e.g., the overall average value) for the texture. The LOD value for each pixel

is utilized to determine an appropriate mip-map level such that the texel corresponds one-to-one with the screen space pixel.

The LOD filtering provided by mip-mapping is configured to prevent moiré interference patterns, aliasing, and rendering artifacts by scaling and filtering a full resolution texture map into multiple lower resolution versions. The integer portion of the computed LOD value is used to select the correct mip-map level corresponding to the primitive's distance from the viewpoint. For example, a 4-bit LOD value may be utilized to access one of 0 through 15 mip-map levels. The fractional portion of the computed LOD value is used to configure a blending operation performed on corresponding texels in corresponding mip-map levels. For example, the fractional portion is used to assign different weights to texels of corresponding mip-map levels.

The shader operations do not utilize LOD parameter values that are outside of the range of zero to $\log_2(N)$. Thus, the raster module 215 is configurable to clamp the LOD value to a maximum value if a computed LOD value is greater than a predetermined range. Similarly, the LOD may be clamped to a minimum value if a computed LOD value is less than the predetermined range. In other words, when the computed LOD parameter corresponds to a texel:pixel ratio, for the primitive, that is larger than the coarsest (smallest) mip-map level, the LOD value may be clamped to the maximum range value. When the computed LOD parameter corresponds to a texel:pixel ratio that is smaller than the finest (largest) mip-map level, the LOD value may be clamped to this minimum range value. In one implementation, once the LOD value is clamped to the minimum or maximum range value, the fractional portion may also be clamped. Alternatively the fractional portion may be discarded. For example, in a case where the maximum LOD parameter value has been exceeded (e.g., mip-map level 15), the lowest resolution version will be fetched by the data fetch module 220. The fractional portion would be insignificant with respect to any blending process.

The gatekeeper module 220 may be communicatively coupled to the raster module 215. The gatekeeper module 220 may also be communicatively coupled to the scoreboard 245. The gatekeeper module 220 provides data flow control (e.g., regulates the flow of pixels through the pipeline). The gatekeeper module 220 attempts to keep the pipeline as full as possible, such that the modules operate on substantially every clock. To facilitate data flow control, the scoreboard 245 tracks the entry and retirement of pixels in the pipeline. In one implementation, a corresponding bit in the scoreboard 245 is set when a pixel packet is received from the raster module 215 by the gatekeeper module 220. The corresponding bit in the scoreboard 245 is reset as the pixel packet drains out of the graphic processor 200.

The scoreboard 245 is utilized by the gatekeeper module 220 for scheduling, load balancing, resource allocation and hazard avoidance of pixel packets. The gatekeeper module 220 also utilizes the scoreboard 245 to ensure cache coherency for the data fetch 225 and data-write modules 235. By tracking the entry and retirement of pixel packet utilizing the scoreboard 245, the gatekeeper module 220 prevents coincident pixels from entering the micro-shader module 230. The gatekeeper module 220 also utilizes the scoreboard 245 to support recirculation of pixel packets for complex shader operations. The gatekeeper module 220 may also detect micro-shader 230 idle conditions. The gatekeeper module 220 may also handle debug register reads.

The data fetch module 225 may be communicatively coupled to the gatekeeper module 220 and a memory controller 265. The data fetch module 225 may also be communicatively coupled to the fetch cache 250. The data fetch module 225 performs color, texture and depth data

reads from memory (not shown) for each pixel packet flowing through the pipeline. The data fetch module 225 reads data from memory and buffers the data in the fetch cache 250. In one implementation, the data fetch module 225 reads data from memory in 128 bit chunks (e.g., the width of the communication channel between the graphics processor and the memory controller). In one implementation the fetch cache 250 may be partitioned into a color cache, texture cache and depth cache. Alternatively the fetch cache 250 may be composed of separate caches for storing color, texture and depth information. The data fetch module 225 inserts the appropriate color, texture and depth information into each pixel packet prior to sending the pixel packet onto the micro-shader module 230.

It is appreciated that because previous and subsequent pixels are spatially related to each other, such that there is a reasonable probability that the color, texture and/or depth data buffered in the fetch cache 250 may be utilized from subsequent pixel packets. Therefore, the bandwidth utilization of the communication channel between the graphics processor 200 and the memory controller 265 may be reduced, if a subsequent pixel packets can utilize color, texture and/or depth data buffered in the fetch cache 250, which were retrieved for a previous pixel packet.

The micro-shader module 230 may be communicatively coupled to the data fetch module 225. The micro-shader module 230, in one implementation, includes one or more arithmetic logic units (ALU). Each arithmetic logic unit implements a unified arithmetic function. In one implementation, each ALU performs a plurality of scalar arithmetic logic operations in the form of $[a*b \text{ "op" } c*d]$ on a set of input parameters, where "op" represents a programmable operation and * is a multiplication operation. The input parameters may be one or more parameters from a given pixel packet, one or more previous result values and/or one or more constants. The arithmetic logic operation may implement functions such as a fog algorithm, alpha blend algorithm, alpha test algorithm and/or a texture combine operation. The ALU may simultaneously perform different operations, thereby improving the functionality of the ALU and decreasing power consumption, as compared to conventional graphics processors. The output of the micro-shader module 230 is a color and depth value for a given pixel of the display.

The data write module 235 may be communicatively coupled to the ALU modules 230 and the memory controller 270. The data write module 235 may also be communicatively coupled to a write buffer 260 and/or the gatekeeper module 220. The data write module 235 send the resultant pixel data, calculated by the ALU module, out to memory controller 265. In one implementation, the communication channel between the graphics processor 200 and the memory controller 270 is 128-bits wide and the pixel data is 16-bits. Therefore, the resultant pixel data is buffered in the write buffer until 128 bits are accumulated. The accumulated resultant pixel data is output to the memory controller 270 for storage in memory. By buffering the resultant pixel data, efficient use of the communication channel bandwidth may be achieved.

The data write module 235 may also indicate retired writes to the scoreboard 245, thereby keeping data consistent. The data write module 235 may also recirculate pixel packets for further processing. In one implementation, the data write module 235 may optionally dither to 565 colors.

Each module 210–235 of the graphics processor 200 is sub-block clocked. Sub-block clocking enables disabling of the clock tree for a given sub-block which is idle (e.g., command and/or data path). Furthermore, the logic is designed to minimize active regions of the logic cone. For example, the logic cone is optimized such that power is

consumed when an operation is performed as opposed to every clock cycle. Thus, if an operation is not performed on a pixel packet in a block of a given module 210–235 during a clock cycle, power savings propagate down the logic tree of the pipeline.

Referring now to FIGS. 3A and 3B, a flow diagram of an exemplary method of deferred transform evaluation, in accordance with one embodiment of the present invention, is shown. As depicted in FIGS. 3A and 3B, in conjunction with FIGS. 4A and 4B, the method of deferred transform evaluation includes receiving vertex data of a geometric primitive from a host interface, at 310. At 315, it is determined by an interlocking mechanism if the new vertex data will overwrite vertex data currently being utilized. If the new vertex data received from the host interface will overwrite vertex data currently involving a draw command, the host interface is held off from loading the new vertex data into the vertex buffer 400, at 320. If the new vertex data will not overwrite vertex data currently involving a draw command, the host interface loads the new vertex data (e.g., 440, 450, 460) into the vertex buffer 400 and clears the an associated transform indicator bit (e.g., 435, 445, 455) in the vertex buffer 400, at 325.

The method of deferred transform evaluation further includes receiving a draw command from the host interface, at 340. At 345, the vertex data corresponding to the draw command is retrieved from the vertex buffer 400. At 350, it is determined, from the state of transform indicator bit, if the retrieved vertex data has been transformed. If the vertex data (e.g., 440) is three-dimensional data of a graphics primitives, as indicated by the corresponding transform indicator bits (e.g., 435), the vertex data is transformed into a two-dimensional representation, at 355. At 360, the transformed vertex data (e.g., 440') is written back to the vertex buffer 400' and the corresponding transform indicator bit (e.g., 435') is set. If the retrieved vertex data (e.g., 430) has already been transformed into a two-dimensional representation, the transform process 355 and corresponding write back process 360 are not performed.

Accordingly, transforms are not automatically performed when vertex data is loaded into the vertex buffer 400. In addition, if a given vertex data has already been transformed as a result of a previous draw command of a primitive that shared the given vertex datum, the repeated transform of the given vertex data may be advantageously eliminated. Eliminating the repeated transform of vertex data currently stored in the vertex buffer 400 reduces power consumed by the graphics processor. Furthermore, it is appreciated that the transform indicator bit enables storage of pre-transform vertex data and post-transform vertex data in a single vertex buffer 400. Hence, the hardware (e.g., transistors implementing the registers utilized to store vertex data) of the graphics processor may be reduced. In addition, the sharing of a single vertex buffer allows the same underlying hardware units (multipliers, adders and the like) to be utilized by both the transformation and setup calculations by virtue of the inherent arbitration between said calculations, further reducing cost and complexity.

Referring now to FIG. 5, a flow diagram of a method of interpolating level-of-detail, in accordance with one embodiment of the present invention, is shown. As depicted in FIG. 5, the method of interpolating the level-of-detail begins with receipt of parameters for each vertex of a given geometric primitive, at 510. Each vertex is specified by a plurality of parameters, such as a primitive X coordinate value, a Y coordinate value, and a W coordinate value (e.g., view distance). At 515, the level-of-detail (LOD) is calculated at each vertex. The LOD of each vertex may be calculated according to Equation 2 for each texture. In one

implementation, the LOD of each vertex varies across the triangle as a function of $1/W^2$.

At 515, the LOD of the pixels of the primitive may be interpolated from the LOD values of the primitive vertices according to Equation 4. The LOD value is between zero and $\text{Log}_2(N)$, where N is the maximum texture dimension supported by the device. The LOD value includes an integer portion and a fractional portion. In one implementation, the LOD value is encoded as an eight-bit representation. Four bits (e.g., integer portion) indicate which of a plurality of mip-map levels to fetch and four bits (e.g., fractional portion) are utilized to interpolate the texel in the specified mip-map level.

At optional process 525, the integer portion and/or the fractional portion of the LOD value may be clamped. For example, if the calculated LOD is less than a zero, than the integer portion of the LOD value is clamped to zero and the fractional portion may be discarded. If the calculated LOD is greater than $\text{Log}_2(N)$, than the integer portion of the LOD value is clamped to maximum resolution and the fractional portion may be discarded. Alternatively, if the calculated LOD is less than a zero, than the integer portion of the LOD value is clamped to zero and the fractional portion may also be clamped. If the calculated LOD is greater than $\text{Log}_2(N)$, than the integer portion of the LOD value is clamped to maximum resolution and the fractional portion may also be clamped. The clamped LOD value may then be output, at 530.

In the conventional art, the LOD is calculated for each pixel in a given primitive based on the pixel's other parameters. In embodiments of the present invention, the level-of-detail (LOD) is calculated once per vertex of a given primitive. It is appreciated that a given triangle has only three vertexes, but typically has hundreds of pixels. Hence, the computational workload of graphics processors, in accordance with embodiments of the present invention, is advantageously reduced as compared to conventional graphics processors because interpolation is more efficient than computation of the LOD values according to the conventional technique. The reduction in computation workload reduces power consumption by the graphics processors.

The foregoing descriptions of specific embodiments of the present invention have been presented for purposes of illustration and description. They are not intended to be exhaustive or to limit the invention to the precise forms disclosed, and obviously many modifications and variations are possible in light of the above teaching. The embodiments were chosen and described in order to best explain the principles of the invention and its practical application, to thereby enable others skilled in the art to best utilize the invention and various embodiments with various modifications as are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the claims appended hereto and their equivalents.

What is claimed is:

1. A method of performing transformations on a geometric primitive comprising:

storing a non-transformed first geometric primitive into a first slot of a vertex buffer, wherein said vertex buffer comprises a plurality of discrete slots, each slot for storing vertex data associated with a particular geometric primitive and a respective transformation indicator bit;

responsive to said storing, clearing a first transformation indicator bit associated with said first slot, wherein a second slot of said vertex buffer stores a transformed second geometric primitive and wherein a second transformation indicator bit associated with said second slot is set;

13

in response to a draw command involving said first non-transformed geometric primitive, transforming said non-transformed first geometric primitive to produce a transformed first geometric primitive;
 storing said transformed first geometric primitive in said first slots wherein said transformed first geometric primitive overwrites said non-transformed geometric primitive; and
 setting said first transformation indicator bit in response to said transforming said non-transformed first graphic primitive.

2. The method as described in claim 1, wherein said storing said non-transformed first geometric primitive is performed under software control via a host interface.

3. The method as described in claim 2, wherein said storing said non-transformed first geometric primitive into said first slot is prevented while a draw command involving said first slot is pending.

4. The method as described in claim 1, wherein said transformation comprises transforming three dimensional vertex position of said non-transformed first geometric primitive to two dimensional vertex position.

5. The method as described in claim 4, wherein said vertex buffer comprises data of transformed geometric primitives and data of non-transformed geometric primitives within said plurality of slots.

6. The method as described in claim 1, wherein said transforming is performed by a setup module of a graphics pipeline and wherein further said setup module is coupled to said vertex buffer and coupled to supply transformed vertex data to a raster module.

7. A vertex buffer for a graphics pipeline, said vertex buffer comprising:

a plurality of slots, each slot for storing vertex data associated with a particular geometric primitive and each slot comprising a respective transformation indicator bit indicating a transformation state of said particular geometric primitive;

wherein a transformation indicator bit associated with a non-transformed geometric primitive is cleared in response to said non-transformed geometric primitive having been loaded into a slot of said vertex buffer under control of a host interface; and

wherein said transformation indicator bit is set in response to said non-transformed geometric primitive having been transformed by a setup module to produce a transformed geometric primitive that is stored in said vertex buffer.

8. The vertex buffer as described in claim 7, further comprising an interlock mechanism preventing said a slot of said vertex buffer from being overwritten while a draw command is pending that involves vertex data stored in said slot.

9. The vertex buffer as described in claim 7, wherein said plurality of slots comprises data of transformed geometric primitives and data of non-transformed geometric primitives.

10. A graphics processor comprising:

a vertex buffer for storing a plurality of vertex data and a plurality of transform indicator bits, wherein each transform indicator bit is associated with a respective one of said vertex data;

a setup module communicatively coupled to said vertex buffer, for receiving a draw command and for transforming given non-transformed vertex data to transformed vertex data in response to said draw command

14

when a corresponding one of said transform indicator bits is in a first state and for saving said transformed vertex data back to said vertex buffer and for setting said one of said transform indicator bits to a second state; and

a raster module communicatively coupled to said setup module for receiving a pixel packet comprising said transformed vertex data and an instruction corresponding to said draw command.

11. The graphics processor as described in claim 10, wherein said given non-transformed vertex data comprises a three-dimensional primitive position.

12. The graphics processor as described in claim 10, wherein said transformed vertex data comprises a two-dimensional primitive position.

13. The graphics processor as described in claim 10, wherein said one of said transform indicator bits is set to said first state in response to said non-transformed vertex data having been loaded into said vertex buffer under control of a host interface.

14. The graphics processor as described in claim 13, wherein said setup module holds off said host interface from loading said non-transformed vertex data into said vertex buffer when a previous vertex data to be overwritten in said vertex buffer is being operated on by a pending draw command.

15. A method of deferring evaluation of a transform comprising:

buffering a plurality of vertex data;

receiving a draw command;

accessing a given vertex data, corresponding to said draw command, and an associated transform indicator bit;

transforming said given vertex data to produce transformed vertex data if said associated transform indicator bit is cleared indicating said given vertex data comprises non-transformed data; and

overwriting said given vertex data with said transformed vertex data and setting said associated transform indicator bit to indicate that said given vertex data has been transformed.

16. The method as described in claim 15, further comprising:

receiving new vertex data; and

overwriting a particular vertex data with said new vertex data and clearing a transform indicator bit associated with said new vertex data.

17. The method as described in claim 16, further comprising:

determining if said particular vertex data is currently being processed as a result of a previous draw command; and

holding off said overwriting said particular vertex data with said new vertex data until said processing of said particular vertex data is completed.

18. The method as described in claim 15, wherein said transforming said given vertex data comprises transforming a three-dimensional primitive position to a two-dimensional primitive position.

19. The vertex buffer as described in claim 7, wherein said transformed geometric primitive is stored in said slot overwriting said non-transformed geometric primitive.

20. The graphics processor as described in claim 10, wherein said transformed vertex data overwrites said non-transformed vertex data.